# Relaxed algorithms for p-adic numbers 

Jérémy Berthomieu, Joris van der Hoeven, Grégoire Lecerf

## To cite this version:

Jérémy Berthomieu, Joris van der Hoeven, Grégoire Lecerf. Relaxed algorithms for p-adic numbers. 2010. hal-00486680

HAL Id: hal-00486680
https://hal.science/hal-00486680
Preprint submitted on 26 May 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Relaxed algorithms for $\boldsymbol{p}$-ADIC NUMBERS* 

Jérémy Berthomieu<br>Laboratoire d'Informatique<br>UMR 7161 CNRS<br>École polytechnique<br>Route de Saclay<br>91128 Palaiseau Cedex<br>France<br>Email: berthomieu@lix.polytechnique.fr<br>Web: www.lix.polytechnique.fr/~berthomieu

Joris van der Hoeven<br>Laboratoire d'Informatique<br>UMR 7161 CNRS<br>École polytechnique<br>Route de Saclay<br>91128 Palaiseau Cedex<br>France<br>Email: vdhoeven@lix.polytechnique.org<br>Web: www.lix.polytechnique.fr/~vdhoeven

Grégoire Lecerf<br>Laboratoire de Mathématiques<br>UMR 8100 CNRS<br>Université de Versailles<br>45 avenue des États-Unis<br>78035 Versailles Cedex<br>France<br>Email: Gregoire.Lecerf@math.uvsq.fr<br>Web: www.math.uvsq.fr/~lecerf

May 26, 2010


#### Abstract

Current implementations of $p$-adic numbers usually rely on so called zealous algorithms, which compute with truncated $p$-adic expansions at a precision that can be specified by the user. In combination with Newton-Hensel type lifting techniques, zealous algorithms can be made very efficient from an asymptotic point of view.

In the similar context of formal power series, another so called lazy technique is also frequently implemented. In this context, a power series is essentially a stream of coefficients, with an effective promise to obtain the next coefficient at every stage. This technique makes it easier to solve implicit equations and also removes the burden of determining appropriate precisions from the user. Unfortunately, naive lazy algorithms are not competitive from the asymptotic complexity point of view. For this reason, a new relaxed approach was proposed by van der Hoeven in the 90 's, which combines the advantages of the lazy approach with the asymptotic efficiency of the zealous approach.

In this paper, we show how to adapt the lazy and relaxed approaches to the context of $p$-adic numbers. We report on our implementation in the C++ library algebramix of Mathemagix, and show significant speedups in the resolution of $p$-adic functional equations when compared to the classical Newton iteration.


KEYWORDS: p-adic numbers, power series, algorithms
A.M.S. SUBJECT CLASSIFICATION: $68 \mathrm{~W} 30,11 \mathrm{Y} 40,11 \mathrm{Y} 16$

[^0]
## 1. INTRODUCTION

Let $R$ be an effective commutative ring, which means that algorithms are available for all the ring operations. Let $(p)$ be a proper principal ideal of $R$. Any element $a$ of the completion $R_{p}$ of $R$ for the $p$-adic valuation can be written, in a non unique way, as a power series $\sum_{i \geqslant 0} a_{i} p^{i}$ with coefficients in $R$. For example, the completion of $\mathbb{K}[x]$ for the ideal $(x)$ is the classical ring of power series $\mathbb{K}[[x]]$, and the completion of $\mathbb{Z}$ for any prime integer $p$ is the ring of $p$-adic integers written $\mathbb{Z}_{p}$.

In general, elements in $R_{p}$ give rise to infinite sequences of coefficients, which cannot be directly stored in a computer. Nevertheless, we can compute with finite but arbitrarily long expansions of $p$-adic numbers. In the so called zealous approach, the precision $n$ of the computations must be known in advance, and fast arithmetic can be used for computations in $R /\left(p^{n}\right)$. In the lazy framework, $p$-adic numbers are really promises, which take a precision $n$ on input, and provide an $n$-th order expansion on output.

In [Hoe97] appeared the idea that the lazy model actually allows for asymptotically fast algorithms as well. Subsequently [Hoe02], this compromise between the zealous and the naive lazy approaches has been called the relaxed model. The main aim of this paper is the design of relaxed algorithms for computing in the completion $R_{p}$. We will show that the known complexity results for power series extend to this setting. For more details on the power series setting, we refer the reader to the introduction of [Hoe02].

### 1.1. Motivation

Completion and deformation techniques come up in many areas of symbolic and analytic computations: polynomial factorization, polynomial or differential system solving, analytic continuation, etc. They make an intensive use of power series and $p$-adic integers.

### 1.1.1. Recursive equations

The major motivation for the relaxed approach is the resolution of algebraic or functional equations. Most of the time, such equations can be rewritten in the form

$$
\begin{equation*}
Y=\Phi(Y), \tag{1}
\end{equation*}
$$

where the indeterminate $Y$ is a vector in $R_{p}^{d}$ and $\Phi$ some algebraic or more complicated expression with the special property that

$$
\tilde{y}-y \in p^{n} R_{p}^{d} \Longrightarrow \Phi(\tilde{y})-\Phi(y) \in p^{n+1} R_{p}^{d}
$$

for all $y, \tilde{y} \in R_{p}^{d}$ and $n \in \mathbb{N}$. In that case, the sequence $0, \Phi(0), \Phi^{2}(0), \ldots$ converges to a solution $y \in R_{p}^{d}$ of (1), and we call (1) a recursive equation.

Using zealous techniques, the resolution of recursive equations can often be done using a Newton iteration, which doubles the precision at every step [BK78]. Although this leads to good asymptotic time complexities in $n$, such Newton iterations require the computation and inversion of the Jacobian of $\Phi$, leading to a non trivial dependence of the asymptotic complexity on the size of $\Phi$ as an expression. For instance, at higher dimensions $d$, the inversion of the Jacobian usually involves a factor $O\left(d^{3}\right)$, whereas $\Phi$ may be of size $O(d)$. We shall report on such examples in Section 5.

The main rationale behind relaxed algorithms is that the resolution of recursive equations just corresponds to a relaxed evaluation of $\Phi$ at the solution itself. In particular, the asymptotic time complexity to compute a solution has a linear dependence on the size of $\Phi$. Of course, the technique does require relaxed implementations for all operations involved in the expression $\Phi$. The essential requirement for a relaxed operation $\varphi$ is that $\varphi(y)_{n}$ should be available as soon as $y_{0}, \ldots, y_{n}$ are known.

### 1.1.2. Elementary operations

A typical implementation of the relaxed approach consists of a library of basic relaxed operations and a function to solve arbitrary recursive equations built up from these operations. The basic operations typically consist of linear operations (such as addition, shifting, derivation, etc.), multiplication and composition. Other elementary operations (such as division, square roots, higher order roots, exponentiation) are easily implemented by solving recursive equations. In several cases, the relaxed approach is not only elegant, but also gives rise to more efficient algorithms for basic operations.

Multiplication is the key operation and Sections 2, 3 and 4 are devoted to it. In situations were relaxed multiplication is as efficient as naive multiplication (e.g. in the naive and Karatsuba models), the relaxed strategy is optimal in the sense that solving a recursive equation is as efficient as verifying the validity of the solution. In the worst case, as we will see in Proposition 6, relaxed multiplication is $O(\log n)$ times more expensive than zealous multiplication modulo $p^{n}$. If $\mathbb{F}_{p}$ contains many $2^{p}$-th roots of unity, then this overhead can be further reduced to $O\left(\mathrm{e}^{2 \sqrt{\log 2 \log \log n}}\right)$ using similar techniques as in [Hoe07a]. In practice, the overhead of relaxed multiplication behaves as a small constant, even though the most efficient algorithms are hard to implement.

In the zealous approach, the division and the square root usually rely on Newton iteration. In small and medium precisions the cost of this iteration turns out to be higher than a direct call to one relaxed multiplication or squaring. This will be illustrated in Section 6: if $p$ is sufficiently large, then our relaxed division outperforms zealous division.

### 1.1.3. User-friendly interface

An important advantage of the relaxed approach is its user-friendliness. Indeed, the relaxed approach automatically takes care of the precision control during all intermediate computations. A central example is the Hensel lifting algorithm used in the factorization of polynomials in $\mathbb{Q}[x]$ : one first chooses a suitable prime number $p$, then computes the factorization in $\mathbb{Z} / p \mathbb{Z}[x]$, lifts this factorization into $\mathbb{Q}_{p}[x]$, and finally one needs to discover how these $p$-adic factors recombine into the ones over $\mathbb{Q}$ (for details see for instance [GG03, Chapter 15]). Theoretically speaking, Mignotte's bound [GG03, Chapter 6] provides us with the maximum size of the coefficients of the irreducible factors, which yields a bound on the precision needed in $\mathbb{Q}_{p}$. Although this bound is sharp in the worst case, it is pessimistic in several particular situations. For instance, if the polynomial is made of small factors, then the factorization can usually be discovered at a small precision. Here the relaxed approach offers a convenient and efficient way to implement adaptive strategies. In fact we have already implemented the polynomial factorization in the relaxed model with success, as we intend to show in detail in a forthcoming paper.

### 1.2. Our contributions

The relaxed computational model was first introduced in [Hoe97] for formal power series, and further improved in [Hoe02, Hoe07a]. In this article, we extend the model to more general completions $R_{p}$. Although our algorithms will be represented for arbitrary rings $R$, we will mainly focus on the case $R=\mathbb{Z}$ when studying their complexities. In Section 2 we first present the relaxed model, and illustrate it on a few easy algorithms: addition, subtraction, and naive multiplications.

In Section 3, we adapt the relaxed product of [Hoe02, Section 4] to $p$-adic numbers. We first present a direct generalization, which relies on products of finite $p$-expansions. Such products can be implemented in various ways but essentially boil down to multiplying polynomials over $R$. We next focus on the case $R=\mathbb{Z}$ and how to take advantage of fast hardware arithmetic on small integers, or efficient libraries for computations with multiple precision integers, such as GMP $[G+91]$. In order to benefit from this kind of fast binary arithmetic, we describe a variant that internally performs conversion between $p$-adic and 2 -adic numbers in an efficient way. We will show that the performance of $p$-adic arithmetic is similar to power series arithmetic over $R /(p)$.

For large precisions, such conversions between $p$-adic and 2 -adic expansions involve an important overhead. In Section 4 we present yet another blockwise relaxed multiplication algorithm, based on the fact that $R_{p} \cong R_{p^{k}}$ for all $k \geqslant 1$. This variant even outperforms power series arithmetic over $R /(p)$. For large block sizes $k$, the performance actually gets very close to the performance of zealous multiplication.

In Section 5, we recall how to use the relaxed approach for the resolution of recursive equations. For small dimensions $d$, it turns out that the relaxed approach is already competitive with more classical algorithm based on Newton iteration. For larger numbers of variables, we observe important speed-ups.

Section 6 is devoted to division. For power series, relaxed division essentially reduces to one relaxed product [Hoe02, Section 3.2.2]. We propose an extension of this result to $p$-adic numbers. For medium precisions, our algorithm turns out to be competitive with Newton's method.

In Section 7, we focus on the extraction of $r$-th roots. We cover the case of power series in small characteristic, and all the situations within $\mathbb{Z}_{p}$. Common transcendental operations such as exponentiation and logarithm are more problematic in the $p$-adic setting than in the power series case, since the formal derivation of $p$-adic numbers has no nice algebraic properties. In this respect, $p$-adic numbers rather behave like floating point numbers. Nevertheless, it is likely that holonomic functions can still be evaluated fast in the relaxed setting, following [Bre76, CC90, Hoe99, Hoe01, Hoe07b]. We also refer to [Kob84, Kat07] for more results about exponentiation and logarithms in $\mathbb{Q}_{p}$.

Algorithms for $p$-adic numbers have been implemented in several libraries and computer algebra systems: P-Pack [Wan84], Maple, Magma, Pari-gp [PAR08], Mathematica [DS04], Sage [S+09], Flint [HH09], etc. These implementations all use the zealous approach and mainly provide fixed-precision algorithms for $R=\mathbb{Z}$. Only Sage also proposes a lazy interface. However, this interface is not relaxed and therefore inefficient for large precisions.

Most of the algorithms presented in this paper have been implemented in the C++ open source library algebramix of Mathemagix [H+02] (revision 4791, freely available from http://www.mathemagix.org). Although we only report on timings for $p$-adic integers, our code provides support for general effective Euclidean domains $R$.

## 2. Data structures and naive implementations

In this section we present the data structures specific to the relaxed approach, and the naive implementations of the ring operations in $R_{p}$.

### 2.1. Finite $\boldsymbol{p}$-adic expansions

As stated in the introduction, any element $a$ of the completion $R_{p}$ of $R$ for the $p$-adic valuation can be written, in a non unique way, as a power series $\sum_{i \geqslant 0} a_{i} p^{i}$ with coefficients in $R$. Now assume that $M$ is a subset of $R$, such that the restriction of the projection map $\pi: R \rightarrow R /(p)$ to $M$ is a bijection between $M$ and $R /(p)$. Then each element $a$ admits a unique power series expansion $\sum_{i \geqslant 0} a_{i} p^{i}$ with $a_{i} \in M$. In the case when $R=\mathbb{Z}$ and $p \in \mathbb{N} \backslash\{0\}$, we will always take $M=\{0, \ldots, p-1\}$.

For our algorithmic purposes, we assume that we are given quotient and remainder functions by $p$

$$
\begin{aligned}
& \text { quo }(\cdot, p): R \rightarrow R \\
& \operatorname{rem}(\cdot, p): R \rightarrow M,
\end{aligned}
$$

so that we have

$$
a=\text { quo }(a, p) p+\operatorname{rem}(a, p),
$$

for all $a \in R$.
Polynomials $\sum_{i=0}^{n-1} a_{i} p^{i} \in M[p]$ will also be called finite $p$-adic expansions at order $n$. In fact, finite $p$-adic expansions can be represented in two ways. On the one hand, they correspond to unique elements in $R$, so we may simply represent them by elements of $R$. However, this representation does not give us direct access to the coefficients $a_{i}$. By default, we will therefore represent finite $p$-adic expansions by polynomials in $M[p]$. Of course, polynomial arithmetic in $M[p]$ is not completely standard due to the presence of carries.

### 2.2. Classical complexities

In order to analyze the costs of our algorithms, we denote by $\mathrm{M}(n)$ the cost for multiplying two univariate polynomials of degree $n$ over an arbitrary ring $A$ with unity, in terms of the number of arithmetic operations in $A$. Similarly, we denote by $\mathrm{I}(n)$ the time needed to multiply two integers of bit-size at most $n$ in the classical binary representation. It is classical [SS71, CK91, Für07] that $\mathrm{M}(n)=O(n \log n \log \log n)$ and $\mathrm{I}(n)=$ $O\left(n \log n 2^{\log ^{*} n}\right)$, where $\log ^{*}$ represents the iterated $\operatorname{logarithm}$ of $n$. Throughout the paper, we will assume that $\mathrm{M}(n) / n$ and $\mathrm{I}(n) / n$ are increasing. We also assume that $\mathrm{M}(O(n))=O(\mathrm{M}(n))$ and $\mathrm{I}(O(n))=O(\mathrm{I}(n))$.

In addition to the above complexities, which are classical, it is natural to introduce $I_{p}(n)$ as the time needed to multiply two $p$-adic expansions in $\mathbb{Z}_{p}$ at order $n$ with coefficients in the usual binary representation. When using Kronecker substitution for multiplying two finite $p$-adic expansions, we have $\mathbf{I}_{p}(n)=\mathrm{I}(n(\log p+\log n))$ [GG03, Corollary 8.27]. We will assume that $\mathrm{I}_{p}(n) / n$ is increasing and that $\mathrm{I}_{p}(O(n))=O\left(\mathrm{I}_{p}(n)\right)$.

It is classical that the above operations can all be performed using linear space. Throughout this paper, we will make this assumption.

### 2.3. The lazy and relaxed computational models

For the description of our relaxed algorithms, we will follow [Hoe02] and use a C ++ style pseudo-code, which is very similar to the actual $\mathrm{C}++$ implementation in Mathemagix. As in [Hoe02], we will not discuss memory management related issues, which have to be treated with care in real implementations, especially when it comes to recursive expansions (see Section 5 below).

The main class Padic $_{p}$ for $p$-adic numbers really consists of a pointer (with reference counting) to the corresponding "abstract representation class" Padic_rep $p$. On the one hand, this representation class contains the computed coefficients $\varphi: M[p]$ of the number $a$ up till a given order $n$ : $\mathbb{N}$. On the other hand, it contains a "purely virtual method" next, which returns the next coefficient $a_{n}$ :

```
class Padic_rep p
    \varphi:M[p]
    n:\mathbb{N}
    virtual next ()
```

Following C++-terminology, the purely virtual function next is only defined in a concrete representation class which derives from Padic_rep $p_{p}$. For instance, to construct a $p$-adic number from an element in $M$, we introduce the type Constant_Padic_rep $p_{p}$ that inherits from Padic_rep $p_{p}$ in this way:

```
class Constant_Padic_rep \({ }_{p} \triangleq\) Padic_rep \(p_{p}\)
    \(c: M\)
    constructor \((\tilde{c}: M)\)
        \(c:=\tilde{c}\)
    method next ()
        if \(n=0\) then return \(c\) else return 0
```

In this piece of code $n$ represents the current precision inherited from Padic_rep ${ }_{p}$. The user visible constructor is given by

$$
\operatorname{padic}(c: M) \rightarrow \operatorname{Padic}_{p}:=\left(\operatorname{Padic}_{p}\right) \text { new Constant_Padic_rep } p(c) .
$$

This constructor creates a new object of type Constant_Padic_rep $p_{p}$ to represent $c$, after which it address can be casted to the abstract type $\operatorname{Padic}_{p}$ of $p$-adic numbers. From now on, for the sake of conciseness, we no longer describe such essentially trivial user level functions anymore, but only the concrete representation classes.

It is convenient to define one more public top-level function for the extraction of the coefficient $a_{k}$, given an instance $a$ of Padic $_{p}$ and a positive integer $k: \mathbb{N}$. This function first checks whether $k$ is smaller than the order $a . n$ of $a$. If so, then $a_{k}=(a . \varphi)_{k}$ is already available. Otherwise, we keep increasing $a . n$ while calling next, until $a_{k}$ will eventually be computed. For more details, we refer to [Hoe02, Section 2]. We will now illustrate our computational model on the basic operations of addition and subtraction.

### 2.4. Addition

The representation class for sums of $p$-adic numbers, written Sum_Padic_rep ${ }_{p}$, is implemented as follows:

```
class Sum_Padic_rep \(p_{p} \triangleq\) Padic_rep \(_{p}\)
    \(a, b:\) Padic \(_{p}\)
    \(\gamma: R\)
    constructor ( \(\tilde{a}:\) Padic \(\left._{p}, \tilde{b}: \operatorname{Padic}_{p}\right)\)
        \(a:=\tilde{a} ; b:=\tilde{b} ; \gamma:=0\)
    method next ()
        \(t:=a_{n}+b_{n}+\gamma\)
        \(\gamma:=\mathrm{quo}(t, p)\)
        return \(\operatorname{rem}(t, p)\)
```

In the case when $R=\mathbb{Z}$, we notice by induction over $n$ that we have $\gamma \in\{0,1\}$, each time that we enter next, since $0 \leqslant a_{n}+b_{n}+\gamma \leqslant 2 p-1$. In that case, it is actually more efficient to avoid the calls to quo and rem and replace the method next by

```
method next ()
    \(t:=a_{n}+b_{n}+\gamma\)
    if \(t<p\) then
        \(\gamma:=0\)
        return \(t\)
    else
        \(\gamma:=1\)
        return \(t-p\)
```

Proposition 1. Given two relaxed $p$-adic integers $a$ and $b$, the sum $a+b$ can be computed up till precision $n$ using $O(n \log p)$ bit-operations.

Proof. Each of the additions $a_{k}+b_{k}+\gamma$ and subsequent reductions modulo $p$ take $O(\log p)$ bit-operations.

### 2.5. Subtraction in $\mathbb{Z}_{p}$

In general, the subtraction is the same as the addition, but for the special case when $R=\mathbb{Z}$, we may use the classical school book method. In our framework, this yields the following implementation:

```
class Sub_Padic_rep \(p \geqslant \operatorname{Padic}_{p}\)
    \(a, b:\) Padic \(_{p}\)
    \(\gamma: R\)
    constructor ( \(\tilde{a}:\) Padic \(_{p}, \tilde{b}:\) Padic \(_{p}\) )
        \(a:=\tilde{a} ; b:=\tilde{b} ; \gamma:=0\)
    method next ()
        \(t:=a_{n}-b_{n}-\gamma\)
        if \(t \geqslant 0\) then
        \(\gamma:=0\)
        return \(t\)
        else
        \(\gamma:=1\)
        return \(t+p\)
```

Proposition 2. Given two relaxed p-adic integers $a$ and $b$, the difference $a-b$ can be computed up till precision $n$ using $O(n \log p)$ bit-operations.

Proof. Each call to the function next $\operatorname{costs} O(\log p)$ bit-operations.

### 2.6. Naive product

Here we consider the school book algorithm: each coefficient ( $a b)_{n}$ is obtained from the sum of all products of the form $a_{k} b_{n-k}$ plus the carry involved by the products of the preceding terms. Carries are larger than for addition, so we have to take them into account carefully. The naive method is implemented in the following way:

```
class Naive_Mul_Padic_rep p}\mp@subsup{\}{\}{}\mp@subsup{\mathrm{ Padic_rep}}{p}{
    a,b: Padic
    \gamma: a vector with entries in R, with indices starting at 0.
    constructor (\tilde{a}:\mp@subsup{Padic}{p}{},\tilde{b}:\mp@subsup{\mathrm{ Padic}}{p}{})
        a:=\tilde{a};b:=\tilde{b};
        Initialize }\gamma\mathrm{ with the empty vector
    method next ()
    Append a zero }\mp@subsup{\gamma}{n}{}=0\mathrm{ at the end of }
    t:=0
    for i from 0 to n do
```

$$
\begin{aligned}
& s:=a_{i} b_{n-i}+\gamma_{i} \\
& t:=t+s \\
& \gamma_{i}:=\operatorname{quo}(t, p) \\
& t:=\operatorname{rem}(t, p) \\
& \text { return } t
\end{aligned}
$$

Proposition 3. Given two relaxed p-adic integers $a$ and $b$, the product $a b$ can be computed up till precision $n$ using $O\left(n^{2} \mathbf{I}(\log p)\right)$ bit-operations.

Proof. We show by induction that, when entering in next to compute the coefficient $(a b)_{n}$, the vector $\gamma$ has size $n$ and entries in $M$. This clearly holds for $n=0$. Assume that the hypothesis is satisfied until a certain value $n \geqslant 0$. When entering next the size of $\gamma$ is increased by 1 , so that it will be $n+1$ at the end. Then, at step $i \in\{0, \ldots, n\}$ of the loop we have $s \leqslant(p-1)^{2}+p-1=p^{2}-p$. Since $t \leqslant p-1$ it follows that $s+t \leqslant p^{2}-1$, whence $\gamma_{i} \leqslant p-1$ on exit. Each of the $O\left(n^{2}\right)$ steps within the loop takes $O(\mathrm{I}(\log p))$ bit-operations, which concludes the proof.

Since hardware divisions are more expensive than multiplications, performing one division at each step of the above loop turns out to be inefficient in practice. Especially when working with hardware integers, it is therefore recommended to accumulate as many terms $a_{i} b_{n-i}$ as possible in $s$ before a division. For instance, if $p$ fits 30 bits and if we use 64 bits hardware integers then we can do a division every 16 terms.

### 2.7. Lifting the power series product

In this subsection we assume that we are given an implementation of relaxed power series over $R$, as described in [Hoe02, Hoe07a]. The representation class is written Series_rep ${ }_{R}$ and the user level class is denoted by Series ${ }_{R}$, in the same way as for $p$-adic numbers. Another way to multiply $p$-adic numbers relies on the relaxed product in $R[[p]]$. This mainly requires a lifting algorithm of $M[[p]]$ into $R[[p]]$ and a projection algorithm of $R[[p]]$ onto $M[[p]]$, The lifting procedure is trivial:

```
class Lift_Series_rep \({ }_{R} \triangleq\) Series_rep \(_{R}\)
    a: Padic \({ }_{p}\)
    constructor ( \(\tilde{a}:\) Padic \(_{p}\) )
        \(a:=\tilde{a}\)
    method next ()
        return \(a_{n}\)
```

Let lift denote the resulting function that converts a $p$-adic number $a$ : Padic $p_{p}$ into a series in Series ${ }_{R}$. The reverse operation, project, is implemented as follows:

```
class Project_Padic_rep \(p\) Padic_rep \(p\)
    \(f:\) Series \(_{R}\)
    \(\gamma: R\)
    constructor \(\left(\tilde{f}:\right.\) Series \(\left._{R}\right)\)
        \(f:=\tilde{f} ; \gamma:=0\)
    method next ()
        \(t:=f_{n}+\gamma\)
        \(\gamma:=\mathrm{quo}(t, p)\)
        return \(\operatorname{rem}(t, p)\)
```

Finally the product $c=a b$ is obtained as project (lift (a) lift $(b))$.
Proposition 4. Given relaxed p-adic integers $a$ and $b$, the product $a b$ can be computed up till precision $n$ using $O(\mathrm{I}(\log p+\log n) \mathrm{M}(n) \log n)$ or $O(\mathrm{I}(n(\log p+\log n)) \log n)$ bitoperations.

Proof. The relaxed product of two power series in size $n$ can be done with $O(\mathrm{M}(n) \log n)$ operations in $R$ by [Hoe02, Theorem 4.1]. In our situation the size of the integers in the product are in $O(\log p+\log n)$. Then, by induction, one can easily verify that the size of the carry $\gamma$ does not exceed $O(\log p+\log n)$ during the final projection step. We are done with the first bound.

The second bound is a consequence of the classical Kronecker substitution: we can multiply two polynomials in $\mathbb{Z}[x]$ of size $n$ and coefficients of bit-size $O(\log p)$ with $O(1(n(\log p+\log n)))$ bit operations [GG03, Corollary 8.27].

This strategy applies in full generality and gives a "softly optimal algorithm". It immediately benefits from any improvements in the power series product. Nevertheless, when $n$ is not much larger than $p$, practical implementations of this method involve a large constant overhead. In the next sections, we will therefore turn our attention to "native" counterparts of the relaxed power series products from [vdH02, Hoe07b].

### 2.8. Timings

We conclude this section with some timings for our $\mathrm{C}++$ implementation of naive multiplication inside the Algebramix package of the Mathemagix system [ $\mathrm{H}+02$ ]. Timings are measured using one core of an Intel Xeon X5450 at 3.0 GHz running Linux and Gmp 5.0.0 [G+91]. As a comparison, we display timings on the same platform, obtained for the Maple 13 package Padic. Precisely, we created two random numbers to precision $n$, did their product via the function evalp, and then asked for the coefficient of order $n / 2$. Notice that timings for small precisions are not very relevant for Maple because of the overhead due to the interpreter.

| $n$ | 8 | 16 | 32 | 64 | 128 | 256 | 512 | 1024 |
| :--- | :---: | :---: | :---: | ---: | ---: | ---: | ---: | ---: |
| Naive_Mul_Padic | $p$ | 2.9 | 3.8 | 8.3 | 22 | 68 | 250 | 920 |
| MAPLE 13 | 240 | 320 | 520 | 1200 | 3500 | 11000 | 38000 | 160000 |

Table 1. Naive products, for $p=536870923$, in microseconds.

## 3. RELAXED PRODUCT

In this section, we extend the relaxed product of [Hoe02, Section 4.3.1] to more general $p$-adic numbers. We also present a special version for $R=\mathbb{Z}$, which uses internal base conversions between base 2 and base $p$, and takes better advantage of the fast arithmetic in Gmp [G+91].

### 3.1. Fast relaxed multiplication algorithm

Let $a$ and $b$ denote the two $p$-adic numbers that we want to multiply, and let $c$ be their product. Let us briefly explain the basic idea behind the speed-up of the new algorithm with respect to naive lazy multiplication.

The first coefficient $c_{0}$ is simply obtained as the remainder of $a_{0} b_{0}$ in the division by $p$. The corresponding quotient is stored as a carry in a variable $\gamma$ similar to the one used in Naive_Mul_Padic_rep ${ }_{p}$. We next obtain $c_{1}$ by computing $a_{0} b_{1}+a_{1} b_{0}+\gamma$ and taking the remainder modulo $p$; the quotient is again stored in $\gamma$. At the next stage, which basically requires the computation of $a_{0} b_{2}+a_{1} b_{1}+a_{2} b_{0}+\gamma$, we do a little bit more than necessary: instead of $a_{1} b_{1}$, we rather compute $\left(a_{1}+a_{2} p\right)\left(b_{1}+b_{2} p\right)$. For $c_{3}$, it then suffices to compute $a_{0} b_{3}$ and $a_{3} b_{0}$ since $a_{1} b_{2}+a_{2} b_{1}$ has already been computed as part of $\left(a_{1}+a_{2} p\right)\left(b_{1}+b_{2} p\right)$. Similarly, in order to obtain $c_{4}$, we only need to compute $a_{0} b_{4}, a_{4} b_{0}, a_{3} b_{1}$ and $a_{1} b_{3}$, since $a_{2} b_{2}$ is already known. In order to anticipate more future computations, instead of computing $a_{3} b_{1}, a_{1} b_{3}$, we compute $\left(a_{1}+a_{2} p\right)\left(b_{3}+b_{4} p\right)$ and $\left(a_{3}+a_{4} p\right)\left(b_{1}+b_{2} p\right)$.

In Figure 1 below, the contribution of each $a_{i} b_{j}$ to the product $c_{i+j}$, corresponds to a small square with coordinates $(i, j)$. Each such square is part of a larger square which corresponds to a product

$$
\left(a_{k}+\cdots+a_{k+2^{q}-1} p^{2^{q}-1}\right)\left(b_{l}+\cdots+b_{l+2^{q}-1} p^{2^{q}-1}\right) .
$$

The number inside the big square indicates the stage $k+l$ at which this product is computed. For instance the products

$$
\begin{aligned}
& \left(a_{3}+a_{4} p+a_{5} p^{2}+a_{6} p^{3}\right)\left(b_{7}+b_{8} p+b_{9} p^{2}+b_{10} p^{3}\right), \\
& \left(a_{7}+a_{8} p+a_{9} p^{2}+a_{10} p^{3}\right)\left(b_{3}+b_{4} p+b_{5} p^{2}+b_{6} p^{3}\right) .
\end{aligned}
$$

correspond to the two $4 \times 4$ squares marked with 10 inside.


Figure 1. Relaxed product.
Given a $p$-adic number $a$, it will be convenient to denote

$$
a_{i \ldots j}=a_{i}+a_{i+1} p+\cdots+a_{j-1} p^{j-1-i} .
$$

For any integer $n$, we also define $l_{n}$ to be the largest integer such that $2^{l_{n}}$ divides $n+2$ if $n+2$ is not a power of 2 . Otherwise, if $n+2=2^{m}$, we let $l_{n}=m-1$. For instance, $l_{0}=$ $0, l_{1}=0, l_{2}=1, l_{3}=0, l_{4}=1$, etc.

We can now describe our fast relaxed product. Recall that $\varphi$ is the finite $p$-expansion inherited from Padic_rep $p_{p}$ that stores the coefficients known to order $n$. In the present situation we also use $\varphi$ for storing the anticipated products.

```
class Relaxed_Mul_Padic_rep \({ }_{p} \triangleq\) Padic_rep \(_{p}\)
    \(a, b:\) Padic \(_{p}\)
    \(\gamma^{a}, \gamma^{b}\) : vectors of vectors over \(R\), with indices starting from 0
    constructor ( \(\tilde{a}:\) Padic \(_{p}, \tilde{b}:\) Padic \(_{p}\) )
        \(a:=\tilde{a}, b:=\tilde{b}\)
        Initialize \(\gamma^{a}\) and \(\gamma^{b}\) with the empty vector
    method next ()
            On entry, \(\gamma^{a}\) and \(\gamma^{b}\) have size \(2 n\); resize them to \(2(n+1)\)
            Initialize \(\gamma_{2 n}^{a}\) and \(\gamma_{2 n}^{b}\) with the zero vector of size \(l_{2 n}+1\)
            Initialize \(\gamma_{2 n+1}^{a}\) and \(\gamma_{2 n+1}^{b}\) with the zero vector of size \(l_{2 n+1}+1\)
            \(t^{a}:=0, t^{b}:=0\)
            for \(q\) from 0 to \(l_{n}\) do
```

$$
\begin{aligned}
& \quad k:=(n+2) / 2^{q} \\
& \quad t^{a}+=\gamma_{n, q}^{a}, t^{b}+=\gamma_{n, q}^{b} \\
& \quad t^{a}+=a_{2^{q}-1 \ldots 2^{q+1}-1} b_{(k-1) 2^{q}-1 \ldots k 2^{q}-1} \\
& \quad \text { if } k=2 \text { then break } \\
& \quad t^{b}+=a_{(k-1) 2^{q}-1 \ldots k 2^{q}-1} b_{2^{q}-1 \ldots 2^{q+1}-1} \\
& s^{a}:=\varphi_{n \ldots n+2^{l_{n+1}}+t^{a}} \\
& \text { Replace } \varphi_{n}, \ldots, \varphi_{n+2^{l_{n+1}-1}} \text { by } s_{0}^{a}, \ldots, s_{2^{l_{n+1}-1}}^{a} \\
& \text { if } n+2 \neq 2^{l_{n}+1} \text { then } \gamma_{n+2^{l_{n+1}, l_{n}}}^{a}:=s_{2^{l_{n+1}}}^{a} \\
& s^{b}:=\varphi_{n \ldots n+2^{l_{n+1}}+t^{b}} \\
& \text { Replace } \varphi_{n}, \ldots, \varphi_{n+2^{l_{n+1}-1}} \text { by } s_{0}^{b}, \ldots, s_{2^{l_{n+1}-1}}^{b} \\
& \text { if } n+2 \neq 2^{l_{n}+1} \text { then } \gamma_{n+2^{l_{n+1}}, l_{n}}^{b}:=s_{2^{l_{n+1}}}^{b} \\
& \text { return } \varphi_{n}
\end{aligned}
$$

Example 5. Let us detail how our algorithm works with the first four steps of the multiplication $c=a b$ with

$$
\begin{aligned}
a & =676=4+5 \times 7+6 \times 7^{2}+7^{3} \\
b & =-1=6+6 \times 7+6 \times 7^{2}+6 \times 7^{3}+O\left(7^{4}\right)
\end{aligned}
$$

Computation of $c_{0}$. Since $l_{0}=l_{1}=0$, the entries $\gamma_{0}^{a}, \gamma_{0}^{b}, \gamma_{1}^{a}, \gamma_{1}^{b}$ are set to (0). In the for loop, $q$ takes the single value 0 , which gives $k=2, t^{a}=a_{0} b_{0}=3+3 \times 7$, and $t^{b}=0$. Then we deduce $s^{a}=3+3 \times 7$, and we set $\varphi_{0}=3, \varphi_{1}=3$. In return we have $c_{0}=3$.
Computation of $c_{1}$. We have $l_{1}=0, l_{2}=1$ and $l_{3}=0$, so that $\gamma_{2}^{a}$ and $\gamma_{2}^{b}$ are initialized with $(0,0)$, while $\gamma_{3}^{a}$ and $\gamma_{3}^{b}$ are initialized with (0). In the for loop, $q$ takes again the single value 0 , and $k$ is set to 3 . We obtain $t^{a}=a_{0} b_{1}=3+3 \times 7$ and $t^{b}=a_{1} b_{0}=2+4 \times 7$. It follows that $s^{a}=6+3 \times 7$, and then that $s^{b}=1+1 \times 7+1 \times 7^{2}$. Finally we set $\varphi_{1}=1$, $\varphi_{2}=1, \gamma_{3,0}^{b}=s_{2}^{b}=1$, and we return $c_{1}=1$.
Computation of $c_{2}$. We have $l_{2}=1, l_{4}=1$ and $l_{5}=0$, so that $\gamma_{4}^{a}$ and $\gamma_{4}^{b}$ are initialized with $(0,0)$ and $\gamma_{5}^{a}$ and $\gamma_{5}^{b}$ with ( 0 ). During the first step of the for loop we have $q=0$, $k=4, t^{a}=a_{0} b_{2}=3+3 \times 7$ and $t^{b}=a_{2} b_{0}=1+5 \times 7$. In the second step we have $q=1$, $k=2$, and we add $a_{1 \ldots 3} b_{1 \ldots 3}=\left(a_{1}+a_{2} \times 7\right)\left(b_{1}+b_{2} \times 7\right)=2+4 \times 7^{2}+6 \times 7^{3}$ to $t^{a}$, its value becomes $5+3 \times 7+4 \times 7^{2}+6 \times 7^{3}$. Then we get $s^{a}=6+3 \times 7+4 \times 7^{2}+6 \times 7^{3}$, and then $s^{b}=2 \times 7+5 \times 7^{2}+6 \times 7^{3}$. Finally we set $\varphi_{2}=0, \varphi_{3}=2, \varphi_{4}=5, \varphi_{5}=6$, and return 0 for $c_{2}$.
Computation of $c_{3}$. We have $l_{3}=0, l_{6}=2$ and $l_{7}=0$, hence $\gamma_{6}^{a}$ and $\gamma_{6}^{b}$ are set to $(0,0,0)$, and $\gamma_{7}^{a}$ and $\gamma_{7}^{b}$ to (0). In the for loop, $q$ takes the single value 0 . We have $k=5, t^{a}=$ $a_{0} b_{3}=3+3 \times 7$ and $t^{b}=\gamma_{3,0}^{b}+a_{3} b_{0}=7$. Then we deduce $s^{a}=5+7+7^{2}$ which yields $\gamma_{5,0}^{a}=1$, and then $s^{b}=5+2 \times 7$. In return we thus obtain $c_{3}=5$.
Proposition 6. Given relaxed $p$-adic integers $a$ and $b$, the product $a b$ can be computed up till precision $n$ using $O\left(1_{p}(n) \log n\right)$ bit-operations. For this computation, the total amount of space needed to store the carries $\gamma^{a}$ and $\gamma^{b}$ does not exceed $O(n)$.

Proof. The proof is essentially the same as for [Hoe02, Theorem 4.1], but the carries require additional attention. We shall prove by induction that all the entries of $\gamma^{a}$ and $\gamma^{b}$ are always in $\{0,1\}$ when entering next for the computation of $\varphi_{n}$. This holds for $n=0$. Assume now that it holds for a certain $n \geqslant 0$. After the first step of the loop, namely for $q=0$, we have $t^{a} \leqslant(p-1)^{2}+1 \leqslant p^{2}-1$. After the second step, when $q=1$, we have $t^{a} \leqslant\left(p^{2}-1\right)^{2}+p^{2}-1+1 \leqslant p^{4}-1$. By induction, it follows that $t^{a} \leqslant p^{2^{q+1}}-1$, at the end of the $q$-th step.

At the end of the for loop, we thus get $t^{a} \leqslant p^{2^{l n+1}}-1$. This implies $s^{a} \leqslant 2 p^{2^{l n+1}}-2$, whence $\gamma_{n+2^{l_{n+1}, l_{n}}}^{a} \in\{0,1\}$. The same holds for superscripts $b$ instead of $a$. Notice that if $n+2 \neq 2^{l_{n}+1}$ then $2^{l_{n}+1} \leqslant n+1$, hence $n+2^{l_{n}+1} \leqslant 2 n+1$. This implies that $\gamma_{n+2^{l n+1, l_{n}}}^{a}$ and $\gamma_{n+2^{l n+1}, l_{n}}^{b}$ are well defined.

If $n+2=2^{l_{n+1}}$, then $s_{2^{n+1}}^{a}=s_{2^{l_{n+1}}}^{b}=0$, since $s^{a}$ and $s^{b}$ are bounded by $\left(p^{n+1}-1\right)^{2} / p^{n}<p^{n+2}$. On the other hand, the map $n \mapsto\left(n+2^{l_{n}+1}, l_{n}\right)$ is injective, so that each entry of $\gamma^{a}$ and $\gamma^{b}$ can be set to 1 at most once. It thus follows that all the carries are carefully stored in the vectors $\gamma^{a}$ and $\gamma^{b}$.

If $n+2=2^{l_{n}} u$, with $u \geqslant 2$, then $n+2^{l_{n}+1}+2=2^{l_{n}}(u+2)$, with $u+2 \geqslant 2$. This implies that, when we arrive at order $n^{\prime}=n+2^{l_{n}+1}$, then the value $l_{n^{\prime}}$ is at least $l_{n}$. Therefore all the carries are effectively taken into account. This proves the correctness of the algorithm.

The cost of the algorithm at order $n$ is

$$
O\left(\sum_{i=0}^{n} \sum_{q=0}^{l_{i}} \mathrm{I}_{p}\left(2^{i}\right)\right)=O\left(\sum_{2^{q} \leqslant n} \frac{n}{2^{q}} \mathrm{I}_{p}\left(2^{q}\right)\right)=O\left(\mathrm{I}_{p}(n) \sum_{2^{q} \leqslant n} 1\right)=O\left(\mathrm{I}_{p}(n) \log n\right),
$$

using our assumption that $\mathrm{I}_{p}(n) / n$ is increasing. Finally,

$$
O\left(\sum_{i=0}^{n}\left(1+l_{i}\right)\right)=O\left(\sum_{i=0}^{n} \sum_{q=0}^{l_{i}} 1\right)=O\left(\sum_{2^{q} \leqslant n} \frac{n}{2^{q}}\right)=O(n)
$$

provides enough space for storing all the carries.
In practice, instead of increasing the sizes of carry vectors by one, we double these sizes, so that the cost of the related memory allocations and copies becomes negligible. The same remark holds for the coefficients stored in $\varphi$.

When multiplying finite $p$-adic expansions using Kronecker substitution, we obtain a cost similar to the one of Proposition 4. Implementing a good product for finite $p$-adic expansions requires some effort, since we cannot directly use binary arithmetic available in the hardware. In the next subsection, we show that minor modifications of our relaxed product allow us to convert back and forth between the binary representation in an efficient manner. Finally, notice that in the case when $R=\mathbb{K}[x]$ and $(p)=(x)$, the carries $\gamma^{a}$ and $\gamma^{b}$ are useless.

### 3.2. Variant with conversion to binary representation

In this subsection, we assume that $R=\mathbb{Z}$ and we adapt the above relaxed product in order to benefit from fast binary arithmetic in available in the processor or Gmp. Informally speaking, we perform conversions between bases $p$ and 2 in a relaxed manner.

```
class Binary_Mul_Padic_rep p}\ Padic_rep p
    a,b: Padic
    \beta}\mp@subsup{\beta}{}{a},\mp@subsup{\beta}{}{b},\mp@subsup{\delta}{}{a},\mp@subsup{\delta}{}{b},\gamma: vectors over \mathbb{Z},\mathrm{ with indices starting from 0.
    constructor (\tilde{a}: Padic}p,\tilde{b}:\mp@subsup{\mathrm{ Padic}}{p}{}
        a:=\tilde{a},b:=\tilde{b}
        Initialize }\mp@subsup{\beta}{}{a},\mp@subsup{\beta}{}{b},\mp@subsup{\delta}{}{a},\mp@subsup{\delta}{}{b},\gamma\mathrm{ with empty vectors
    method next ()
        If }n+2\mathrm{ is a power of 2, then
```

Resize $\beta^{a}, \beta^{b}, \delta^{a}, \delta^{b}$, and $\gamma$ to $l_{n}+1$
Fill the new entries with zeros

$$
\begin{aligned}
& \varepsilon^{a}:=a_{n}, \varepsilon^{b}:=b_{n}, \tau:=0 \\
& \text { for } q \text { from } 0 \text { to } l_{n} \text { do } \\
& \text { if } q>0 \text { then } \\
& \quad \varepsilon^{a}:=\beta_{q-1}^{a}+p^{2^{q-1}} \varepsilon^{a} \\
& \quad \varepsilon^{b}:=\beta_{q-1}^{b}+p^{2^{q-1}} \varepsilon^{b} \\
& \text { if } n+2=2^{q+1} \text { then } \\
& \quad \delta_{q}^{a}:=\varepsilon^{a}, \delta_{q}^{b}:=\varepsilon^{b} \\
& \quad \tau+=\delta_{q}^{a} \varepsilon^{b}+\gamma_{q} \\
& \quad \text { if } n+2=2^{q+1} \text { then break } \\
& \quad \tau+=\varepsilon^{a} \delta_{q}^{b} \\
& \beta_{l_{n}}^{a}:=\varepsilon^{a}, \beta_{l_{n}}^{b}:=\varepsilon^{b} \\
& \text { for } q \text { from } l_{n} \text { down to } 0 \text { do } \\
& \quad \gamma_{q}:=q u o\left(\tau, p^{2^{q}}\right), \tau:=\text { rem }\left(\tau, p^{2^{q}}\right) \\
& \text { return } \tau
\end{aligned}
$$

Proposition 7. Given two relaxed p-adic integers $a$ and $b$, the computation of the product a b up till precision $n$ can be done using $O(\mathrm{I}(n \log p) \log n)$ bit-operations and $O(n \log p)$ bit-space.

Proof. When computing $\varphi_{n}$, the vectors $\beta^{a}, \beta^{b}, \delta^{a}, \delta^{b}$, and $\gamma$ are resized to $r_{n}$, where $r_{n}$ is the largest integer such that $2^{r_{n}} \leqslant n+2$. From $2^{r_{n}+1}>n+2 \geqslant 2^{l_{n}+1}$, we deduce that $r_{n}>l_{n}$, which means that the read and write operations in these vectors are licit.

For any integers $n$ and $q$ such that $2^{q+1}<n+2$, we write $\mu(n, q)$ for the largest integer less than $n$ such that $l_{\mu(n, q)}=q$. We shall prove by induction that, when entering next for computing $\varphi_{n}$, the following properties hold for all $q \geqslant 0$ such that $2^{q+1}<n+2$ :
a) $\beta_{q}^{a}=a_{(k-1) 2^{q}-1 \ldots k 2^{q}-1}=a_{\mu(n, q)-2^{q}+1 \ldots \mu(n, q)+1}$ where $k=\frac{\mu(n, q)+2}{2^{q}}$, and similarly for $\beta_{q}^{b}$,
b) $\delta_{q}^{a}=a_{2^{q}-1 \ldots 2^{q+1}-1}$, and similarly for $\delta_{q}^{b}$, and
c) $\gamma_{q} \leqslant 2 p^{2^{q}}$.

These properties trivially hold for when $n=0$. Let us assume that they hold for a certain $n \geqslant 0$.

Now we claim that, at the end of step $q$ of the first loop, the value of $\varepsilon_{n}^{a}$ is $a_{(k-1) 2^{q}-1 \ldots k 2^{q}-1}=a_{n-2^{q}+1 \ldots n+1}$ with $k=(n+2) / 2^{q}$. This clearly holds for when $q=0$ because $\varepsilon^{a}=a_{n}$ and $k=n+2$. Now assume that this claim holds until step $q-1$ for some $q \geqslant 1$. When entering step $q$, we have that $\mu(n, q-1)=n-2^{q-1}$, and part (a) of the induction hypothesis gives us that $\beta_{q-1}^{a}=a_{n-2^{q}+1 \ldots n-2^{q-1}+1}$. From these quantities, we deduce:

$$
\begin{aligned}
\beta_{q-1}^{a}+p^{2^{q-1}} \varepsilon^{a} & =a_{n-2^{q}+1 \ldots n-2^{q-1}+1}+p^{2^{q-1}} a_{n-2^{q-1}+1 \ldots n+1} \\
& =a_{n-2^{q}+1 \ldots n+1}=a_{(k-1) 2^{q}-1 \ldots k 2^{q}-1}
\end{aligned}
$$

with $k=(n+2) / 2^{q}$, which concludes the claim by induction. If $n+2$ is not a power of 2 then part (a) is clearly ensured at the end of the computation of $\varphi_{n}$. Otherwise $n+2=$ $2^{l_{n}+1}$, and $\beta_{l_{n}}^{a}$ is set to $a_{n-2^{l_{n}}+1 \ldots n+1}$, and part (a) is again satisfied when entering the computation of $\varphi_{n+1}$.

When $\delta_{q}^{a}$ is set to $\varepsilon^{a}$, the value of $\varepsilon^{a}$ is $a_{(k-1) 2^{q}-1 \ldots k 2^{q}-1}$ with $k=2$. This ensures that part (b) holds when entering the computation of $\varphi_{n+1}$.

As to (c), during step $q$ of the first loop, the value of $\tau$ is incremented by at most

$$
2\left(p^{2^{q}}-1\right)^{2}+2{p^{2 q}}_{2} \leqslant 2{p^{2 q+1}}^{2 q} 2 p^{2^{q}}+2 .
$$

At the end of this loop, we thus have

$$
\tau \leqslant 2 p^{2^{l_{n+1}}}-2 p+2\left(l_{n}+1\right) .
$$

It follows that $\tau / p^{2^{l_{n}}}<2 p^{2^{l_{n}}}+2\left(l_{n}+1\right) / p^{2^{l_{n}}}$. If $l_{n} \in\{0,1\}$ then it is clear that $2\left(l_{n}+\right.$ $1) \leqslant p^{2^{l n}}$, since $p \geqslant 2$. If $l_{n} \geqslant 1$ then $\frac{d\left(2^{2^{l n}}\right)}{d l_{n}}=\log (2) \log (p) 2^{l_{n}} 2^{2^{l_{n}}} \geqslant 8(\log 2)^{2} \geqslant 2$. We deduce that $2\left(l_{n}+1\right) \leqslant p^{2^{l_{n}}}$ holds for all integer $l_{n} \geqslant 0$. Before exiting the function we therefore have that $\gamma_{l_{n}} \leqslant 2 p^{2^{l n}}, \gamma_{l_{n}-1} \leqslant p^{2^{l_{n-1}}} \leqslant 2 p^{2^{l_{n-1}}}$, etc, which completes the induction.

Since $n+2=2^{l_{n}} u$, with $u \geqslant 2$, we have $n+2^{k}+2=2^{k}\left(2^{l_{n}-k} u+1\right)$, whence $l_{n+2^{k}} \geqslant k$, for any $k \leqslant l_{n}$. All the carries stored in $\gamma$ are therefore properly taken into account. This proves the correctness of the algorithm.

At precision $n$, summing the costs of all the calls to next amounts to

$$
\begin{aligned}
O\left(\sum_{i=0}^{n} \sum_{q=0}^{l_{i}} \mathrm{I}\left(2^{q} \log p\right)\right) & =O\left(\sum_{2^{q} \leqslant n} \frac{n}{2^{q}} \mathrm{I}\left(2^{q} \log p\right)\right) \\
& =O\left(\mathrm{I}(n \log p) \sum_{2^{q} \leqslant n} 1\right) \\
& =O(\mathbf{I}(n \log p) \log n) .
\end{aligned}
$$

Furthermore,

$$
O\left(\sum_{2^{q+1} \leqslant n+2} 2^{q} \log p\right)=O(n \log p)
$$

provides a bound for the total bit-size of the auxiliary vectors $\beta^{a}, \beta^{b}, \delta^{a}, \delta^{b}$, and $\gamma$.
Again, in practice, one should double the allocated sizes of the auxiliary vectors each time needed so that the cost of the related memory allocations and copies becomes negligible. In addition, for efficiency, one should precompute the powers of $p$.

### 3.3. Timings

In following Table 2 , we compare timings for power series over $\mathbb{F}_{p}$, and for $p$-adic integers via Series_Mul_Padic_rep $p_{p}$ of Section 2.7, and via Binary_Mul_Padic_rep $p_{p}$ of Proposition 7. In Series_Mul_Padic_rep ${ }_{p}$ the internal series product is the relaxed one reported in the first line.

| $n$ | 8 | 16 | 32 | 64 | 128 | 256 | 512 | 1024 | 2048 | 4096 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Relaxed mul. in $\mathbb{F}_{p}[[x]]$ | 2 | 7 | 20 | 50 | 140 | 360 | 930 | 2300 | 5700 | 14000 |
| Series_Mul_Padic ${ }_{p}$ | 19 | 59 | 160 | 420 | 1100 | 2600 | 6200 | 14000 | 34000 | 79000 |
| Binary_Mul_Padic ${ }_{p}$ | 8 | 16 | 37 | 89 | 170 | 360 | 800 | 1900 | 4300 | 10000 |
| Naive_Mul_Padic ${ }_{p}$ | 2.9 | 3.8 | 8.3 | 22 | 68 | 250 | 920 | 3600 | 14000 | 56000 |

Table 2. Fast relaxed products, and naive lazy product, for $p=536870923$, in microseconds.

For convenience, we recall the timings for the naive algorithm of Section 2.6 in the last line of Table 2. We see that our Binary_Mul_Padic_rep $p_{p}$ is faster from size 512 on. Since handling small numbers with Gmp is expensive, we also observe some overhead for small sizes, compared to the fast relaxed product of formal power series. On the other hand, since the relaxed product for power series makes internal use of Kronecker substitution, it involves integers that are twice as large as those in Binary_Mul_Padic_rep $p_{p}$. Notice finally that the lifting strategy Series_Mul_Padic_rep ${ }_{p}$, described in Section 2.7, is easy to implement, but not competitive.

## 4. Blockwise product

As detailed in Section 4.1 below for $R=\mathbb{Z}$, the relaxed arithmetic is slower than direct computations modulo $p^{n}$ in binary representation. In [Hoe07a], an alternative approach for relaxed power series multiplication was proposed, which relies on grouping blocks of $k$ coefficients and reducing a relaxed multiplication at order $n$ to a relaxed multiplication at order $n / k$, with FFT-ed coefficients in $M^{2 k-1}$.

Unfortunately, we expect that direct application of this strategy to our case gives rise to a large overhead. Instead, we will now introduce a variant, where the blocks of size $k$ are rather rewritten into an integer modulo $p^{k}$. This aims at decreasing the overhead involved by the control instructions when handling objects of small sizes, and also improving the performance in terms of memory management by choosing blocks well suited to the sizes of the successive cache levels of the platform being used.

We shall start with comparisons between the relaxed and zealous approaches. Then we develop a supplementary strategy for a continuous transition between the zealous and the relaxed models.

### 4.1. Relaxed versus zealous

The first line of Table 3 below displays the time needed for the product modulo $p^{n}$ of two integers taken at random in the range $0, \ldots, p^{n-1}$. The next line concerns the performance of our function binary that converts a finite $p$-expansion of size $n$ into its binary representation. The reverse function, reported in the last line, and written expansion, takes an integer in $0, \ldots, p^{n}-1$ in base 2 and returns its $p$-adic expansion.

| $n$ | 8 | 16 | 32 | 64 | 128 | 256 | 512 | 1024 | 2048 | 4096 | 8192 |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | ---: | ---: | ---: | ---: | :--- |
| mod. mul. | 0.38 | 0.52 | 1.0 | 2.9 | 9.0 | 27 | 85 | 250 | 690 | 1800 | 4500 |
| binary | 1.0 | 2.2 | 4.5 | 9.8 | 20 | 44 | 100 | 250 | 560 | 1400 | 3300 |
| expansion | 2.5 | 5.2 | 10 | 22 | 46 | 96 | 220 | 490 | 1200 | 3000 | 7300 |

Table 3. Zealous product modulo $p^{n}$ and conversions, for $p=536870923$, in microseconds.
Let us briefly recall that binary can be computed fast by applying the classical divide and conquer paradigm as follows:

$$
\operatorname{binary}\left(a, p^{n}\right)=\operatorname{binary}\left(a_{0 \ldots h}, p^{h}\right)+p^{h} \operatorname{binary}\left(a_{h \ldots n}, p^{n-h}\right), \text { where } h=\lfloor n / 2\rfloor,
$$

which yields a cost in $O(\mathrm{I}(n \log p) \log n)$. Likewise, the same complexity bound holds for expansion. Within our implementation we have observed that these asymptotically fast algorithms outperform the naive ones whenever $p^{n}$ is more than around 32 machine words.

Compared to Tables 1 and 2, these timings confirm the theoretical bounds: the relaxed product does not compete with a direct modular computation in binary representation. This is partly due to the extra $O(\log n)$ factor for large sizes. But another reason is the overhead involved by the use of Gmp routines with integers of a few words. In Table 4, we report on the naive and the relaxed products in base $p^{32}$. Now we see that our naive product becomes of the same order of efficiency as the zealous approach up to precision 1024. The relaxed approach starts to win when the precision reaches 256 in base $p^{32}$.

| $k l$ | 32 | 64 | 128 | 256 | 512 | 1024 | 2048 | 4096 | 8192 |
| :--- | :---: | :---: | :---: | :---: | ---: | :---: | :---: | :---: | :---: |
| Naive_Mul_Padic $_{p^{k}}$ | 1.8 | 4.1 | 10 | 27 | 84 | 280 | 1000 | 3900 | 15000 |
| ${\text { Binary_Mul_Padic } p^{k}}^{2}$ | 3.2 | 6.1 | 16 | 53 | 170 | 570 | 1800 | 5300 | 15000 |

Table 4. Relaxed product modulo $\left(p^{k}\right)^{l}$, for $p=536870923$ and $k=32$, in microseconds.

### 4.2. Monoblock strategy

If one wants to compute the product $a b$ of two $p$-adic numbers $a$ and $b$, then: one can start by converting both of them into $p^{k}$-adic numbers $A$ and $B$ respectively, multiply $A$ and $B$ as $p^{k}$-adic numbers, and finally convert $A B$ back into a $p$-adic number. The transformations between $p$-adic and $p^{k}$-adic numbers can be easily implemented:

```
class To_Blocks \({ }_{p^{k}} \triangleq\) Padic_rep \(p_{p^{k}}\)
    a: Padic \(_{p}\)
    constructor ( \(\tilde{a}:\) Padic \(_{p}\) )
        \(a:=\tilde{a}\)
    method next ()
        return binary \(\left(a_{n k \ldots(. .}(n+1) k\right)\)
class From_Blocks \({ }_{p} \triangleq\) Padic_rep \(_{p}\)
    \(a\) : Padic \(_{p^{k}}\)
    \(b: p\)-expansion of size \(k\)
    constructor ( \(\tilde{a}:\) Padic \(_{p^{k}}\) )
        \(a:=\tilde{a}\)
    method next ()
        if \(n \bmod k=0\) then \(b=\mathrm{p} \_\)expansion \(\left(a_{n / k}, p\right)\)
        return \(b_{n \bmod k}\)
```

If to_blocks and from_blocks represent the top level functions then the product $c$ of $a$ and $b$ can be simply obtained as $c=$ from_blocks (to_blocks (a) to_blocks (b)). We call this way of computing products the monoblock strategy.

Notice that choosing $k$ very large is similar to zealous computations. This monoblock strategy can thus be seen as a mix of the zealous and the relaxed approaches. However, it is only relaxed for $p^{k}$-expansions, not for $p$-expansions. Indeed, let $A$ and $B$ still denote the respective $p^{k}$-adic representations of $a$ and $b$, so that $c=$ from_blocks $(C)$, for $C=A B$. Then the computation of $c_{0}$ requires the knowledge of $C_{0}=\overline{A_{0}} B_{0}$, whence it depends on the coefficients $a_{k-1}$ and $b_{k-1}$.

### 4.3. Relaxed blockwise product

We are now to present a relaxed $p$-adic blockwise product. This product depends on two integer parameters $m$ and $k$. The latter still stands for the size of the blocks to be used, while the former is a threshold: below precision $m$ one calls a given product on $p$-expansions, while in large precision an other product is used on $p^{k}$-expansions.

If $a$ and $b$ are the two numbers in $R_{p}$ that we want to multiply as a $p$-expansions, then we first rewrite them $a=a_{0 \ldots m}+p^{m} \bar{a}$ and $b=b_{0 \ldots m}+p^{m} \bar{b}$, where

$$
\bar{a}=a / p^{m}=\sum_{i=0}^{\infty} a_{m+i} p^{i} \quad \text { and } \quad \bar{b}=b / p^{m}=\sum_{i=0}^{\infty} b_{m+i} p^{i} .
$$

Now multiplying $a$ and $b$ gives

$$
c=a_{0 \ldots m} b_{0 \ldots m}+p^{m}\left(a_{0 \ldots m} \bar{b}+\bar{a} b_{0 \ldots m}\right)+p^{2 m} \bar{a} \bar{b},
$$

where the product $\bar{a} \bar{b}$ can be computed in base $p^{k}$, as it is detailed in the following implementation:

```
class Blocks_Mul_Padic_rep \(p\) Padic_rep \(_{p}\)
    \(a, b, c, \bar{a}, \bar{b}:\) Padic \(_{p}\)
    \(\bar{A}, \bar{B}:\) Padic \(_{p^{k}}\)
    constructor ( \(\left.\tilde{a}: \operatorname{Padic}_{p}, \tilde{b}: \operatorname{Padic}_{p}\right)\)
        \(a:=\tilde{a}, b:=\tilde{b}\)
        \(\bar{a}:=a / p^{m}, \bar{b}:=b / p^{m}\)
        \(\bar{A}:=\) to_blocks \((\bar{a}), \bar{B}:=\) to_blocks \((\bar{b})\)
        \(\bar{c}:=\) from_blocks \((\bar{A} \bar{B})\)
        \(c:=a_{0 \ldots m} \bar{b}_{0 \ldots m}+p^{m}\left(a_{0 \ldots m} \bar{b}+\bar{a} b_{0 \ldots m}\right)+p^{2 m} \bar{c}\)
    method next ()
        return \(c_{n}\)
```

In Figure 2 below, we illustrate the contribution of each $a_{i} b_{j}$ to the product $c_{i+j}$ computed with the present blockwise version. In both bases $p$ and $p^{k}$ the naive product is used, and the numbers inside the squares indicate the degrees at which the corresponding product is actually computed.

|  | 11 | 12 |  | 10 |  |  |  | 14 |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 9 | 10 | 11 |  |  |  |  |  |  |  |  |
| 8 | 9 | 10 |  |  |  |  |  |  |  |  |
| 7 | 8 | 9 |  |  |  |  |  |  |  |  |
| 6 | 7 | 8 |  | 6 |  |  |  |  | 10 |  |
| 5 | 6 | 7 |  |  |  |  |  |  |  |  |
| 4 | 5 | 6 |  |  |  |  |  |  |  |  |
| 3 | 4 | 5 |  |  |  |  |  |  |  |  |
| 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |  | 011 |  |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |  |
| 0 | 1 | 2 | 3 |  | 5 | 6 | 7 | 8 | 9 | 10 |

Figure 2. Blockwise product for $m=3$ and $k=4$.
Proposition 8. If $m \geqslant k-1$, then Blocks_Mul_Padic_rep ${ }_{p}$ is relaxed for base $p$.
Proof. It is sufficient to show that the computation of $\bar{c}_{n-2 m}$ only involves terms in $a$ and $b$ of degree at most $n$. In fact $\bar{c}_{n-2 m}$ requires the knowledge of the coefficients of $\bar{A}$ and $\bar{B}$ to degree at most $l=\lfloor(n-2 m) / k\rfloor$, hence the knowledge of the coefficients of $a$ and $b$ to degree $k(l+1)-1+m \leqslant n-2 m+k-1+m=n+k-1-m$, which concludes the proof thanks to the assumption on $m$. Notice that the latter inequality is an equality whenever $n-2 m$ is a multiple of $k$. Therefore $m \geqslant k-1$ is necessary to ensure the product to be relaxed.

### 4.4. Timings

In the following table, we use blocks of size $k=32$, and compare the blockwise versions of the naive product of Section 2.6 to the relaxed one of Section 3.2. The first line concerns the monoblock strategy: below precision 32 we directly use the naive $p$-adic product; for larger precisions we use the naive $p^{k}$-adic product. The second line is the same as the first one except that we use a relaxed $p^{k}$-adic product. In the third line the relaxed blockwise version is used with $m=32$ : we use the naive product for both $p$ - and $p^{k}$-adic expansions. The fourth line is similar except that the fast relaxed product is used beyond precision 32 .

| $n$ | 8 | 16 | 32 | 64 | 128 | 256 | 512 | 1024 | 2048 | 4096 |
| :--- | :---: | :---: | :---: | :---: | ---: | ---: | ---: | ---: | ---: | ---: |
| mono Naive_Mul_Padic $_{p}$ | 3.5 | 5.5 | 11 | 53 | 95 | 190 | 380 | 870 | 2200 | 6500 |
| mono Binary_Mul_Padic $_{p}$ | 3.4 | 5.5 | 11 | 57 | 110 | 220 | 500 | 1200 | 3000 | 7800 |
| blocks Naive_Mul_Padic $_{p}$ | 8.0 | 11 | 17 | 46 | 140 | 320 | 700 | 1600 | 3700 | 9400 |
| blocks Binary_Mul_Padic $_{p}$ | 8.0 | 11 | 17 | 46 | 140 | 330 | 750 | 1700 | 3900 | 9100 |

Table 5. Blockwise products to precision $n$, for $p=536870923$ and $k=32$, in microseconds.

When compared to Table 4, we can see that most of the time within the monoblock strategy is actually spent on base conversions. In fact, the strategy does not bring a significant speed-up for a single product, but for more complex computations, the base conversions can often be factored.

For instance, assume that $a$ and $b$ are two $d \times d$ matrices with entries in $\mathbb{Z}_{p}$. Then the multiplication $c=a b$ involves only $O\left(d^{2}\right)$ base conversions and $O\left(d^{3}\right) p^{k}$-adic products. For large $d$, the conversions thus become inexpensive. In Section 7, we will encounter a similar application to multiple root extraction.

## 5. Application to recursive $p$-Adic numbers

A major motivation behind the relaxed computational model is the efficient expansion of $p$-adic numbers that are solutions to recursive equations. This section is an extension of [Hoe02, Section 4.5] to $p$-adic numbers.

Let us slightly generalize the notion of a recursive equation, which was first defined in the introduction, so as to accommodate for initial conditions. Consider a functional equation

$$
\begin{equation*}
Y=\Phi(Y) \tag{2}
\end{equation*}
$$

where $Y$ is a vector of $d$ unknowns in $R_{p}$. Assume that there exist a $k \in \mathbb{N}^{*}$ and initial conditions $c_{0}, \ldots, c_{k-1} \in M^{d}$, such that for all $n \geqslant k$ and $y, \tilde{y} \in R_{p}^{d}$ with $y_{0}=c_{0}, \ldots$, $y_{k-1}=c_{k-1}$, we have

$$
\begin{equation*}
\tilde{y}-y \in p^{n} R_{p}^{d} \Longrightarrow \Phi(\tilde{y})-\Phi(y) \in p^{n+1} R_{p}^{d} \tag{3}
\end{equation*}
$$

Stated otherwise, this condition means that each coefficient $b_{n}=\Phi(b)_{n}$ with $n \geqslant k$ only depends on previous coefficients $b_{0}, \ldots, b_{n-1}$. Therefore, setting $c=c_{0}+c_{1} p+\cdots+$ $c_{k-1} p^{k-1}$, the sequence $c, \Phi(c), \Phi(\Phi(c))$ converges to a unique solution $b \in R_{p}^{d}$ of (2) with $b_{0}=c_{0}, \ldots, b_{k-1}=c_{k-1}$. We will call (2) a recursive equation and the entries of the solution $b$ recursive $p$-adic numbers.

### 5.1. Implementation

Since $p$ induces an element $\boldsymbol{p}=(p, \ldots, p)$ in $R^{d}$ and an isomorphism $R_{p}^{d} \cong\left(R^{d}\right)_{\boldsymbol{p}}$, we may reinterpret a solution $b=\Phi(b)$ as a $p$-adic number over $R^{d}$. Using this trick, we may assume without loss of generality that $d=1$. In our implementation, recursive numbers are instances of the following class that stores the initial conditions $b_{0}, \ldots, b_{k-1}$ and the equation $\Phi$ :

```
class Recursive_Padic_rep \({ }_{p} \triangleq\) Padic_rep \(_{p}\)
    \(\Phi\) : function from \(R_{p}\) to \(R_{p}\)
    \(b_{0}, \ldots, b_{k-1}\) : initial conditions in \(M\)
    constructor ( \(\tilde{\Phi}\) : function, \(\tilde{b}_{0}, \ldots, \tilde{b}_{k-1}: M\) )
        \(\Phi:=\tilde{\Phi}, b_{0}:=\tilde{b_{0}}, \ldots, b_{k-1}:=\tilde{b_{k-1}}\)
    method next ()
        If \(n<k\) then return \(b_{n}\)
        return \(\Phi(\text { this })_{n}\)
```

In the last line, the expression $\Phi$ (this) means the evaluation of $\Phi$ at the concrete instance of the $p$-adic $b=\Phi(b)$ being currently defined.

Example 9. Consider $\Phi(b)=p b+1$, with one initial condition $b_{0}=1$. It is clear that $b$ is recursive, since the $n$ first terms of $p b$ can be computed from the only $n-1$ first terms of $b$. We have $b_{1}=b_{2}=\cdots=1$. In fact, $b=1 /(1-p)$.

### 5.2. Complexity analysis

If $\Phi$ is an expression built from $L$ constants, sums, differences, and products (all of arity two), then the computation of $b$ simply consists in performing these $L$ operations in the relaxed model. For instance, when using the relaxed product of Proposition 6, this amounts to $O\left(\left.L\right|_{p}(n) \log n\right)$ operations to obtain the $n$ first terms of $b$.

This complexity bound is to be compared to the classical approach via the Newton operator. In fact, one can compute $b$ with fixed-point $p$-adic arithmetic by evaluating the following operator $N_{\Phi}(z)=z-(z-\Phi(z)) /\left(1-\Phi^{\prime}(z)\right)$. There are several cases where the relaxed approach is faster than the Newton operator:

1. The constant hidden behind the " $O$ " of the Newton iteration is higher than the one with the relaxed approach. For instance, if $b$ is really a vector in $R_{p}^{d}$, then the Newton operator involves the inversion of a $d \times d$ matrix at precision $n / 2$, which gives rise to a factor $O\left(d^{3}\right)$ in the complexity (assuming the naive matrix product is used). The total cost of the Newton operator to precision $n$ in $\mathbb{Z}_{p}$ is thus in $O\left(\left(d L+d^{3}\right) I_{p}(n)\right)$. Here $O(d L)$ bounds the number of operations needed to evaluate the Jacobian matrix. In this situation, if $L \ll d^{2}$, and unless $n$ is very large, the relaxed approach is faster. This will be actually illustrated in the next subsection.
2. Even in the case $d=1$, the " $O$ " hides a non trivial constant factor due to a certain amount of "recomputations". For moderate sizes, when polynomial multiplication uses Karatsuba's algorithm, or the Toom-Cook method, the cost of relaxed multiplication also drops to a constant times the cost of zealous multiplication [Hoe02, Hoe07b]. In such cases, the relaxed method often becomes more efficient. This will be illustrated in Section 6 for the division.
3. When using the blockwise method from Section 4 or [Hoe07b] for power series, the overhead of relaxed multiplication can often be further reduced. In practice, we could observe that this makes it possible to outperform Newton's method even for very large sizes.

For more general functional equations, where $\Phi$ involves non-algebraic operations, it should also be noticed that suitable Newton operators $\Phi$ are not necessarily available. For instance, if the mere definition of $\Phi$ involves $p$-expansions, then the Newton operator may be not defined anymore, or one needs to explicitly compute with $p$-expansions. This occurs for instance for $R=\mathbb{Z}$, when $\Phi$ involves the "symbolic derivation" $b \mapsto$ $\sum_{n \geqslant 1} n b_{n} p^{n-1}$.

### 5.3. Timings

In order to illustrate the performance of the relaxed model with respect to Newton iteration, we consider the following family of systems of $p$-adic integers:

$$
\Phi_{d, i}\left(x_{1}, \ldots, x_{d}\right)=1+p \sum_{k=1}^{d}(k+i) x_{k}^{(k+i) \bmod 3}, \quad \text { for } i \in\{1, \ldots, d\}
$$

The number of $p$-adic products grows linearly with $d$. Yet, the total number of operations grows with $d^{2}$.

In Table 6, we compute the 256 first terms of the solution $b=\Phi(b)$ with the initial condition $b=(1, \ldots, 1)+O(p)$. We use the naive product of Section 2.6 and compare to the Newton iteration directly implemented on the top of the routines of Gmp. In fact the time we provide in the line "Matrix multiplication" does not corresponds to a complete implementation of the iteration but only to two products of two $d \times d$ matrices with integers modulo $p^{64}$. These two operations necessarily occurs for inverting the Jacobian matrix to precision $p^{128}$ when using the classical algorithm as described in [GG03, Algorithm 9.2]. This can be seen as a lower bound for any implementation of the Newton method. However the line "Newton implementation" corresponds to our implementation of this method, hence this is an upper bound.

| $d$ | 1 | 2 | 4 | 8 | 16 | 32 | 64 | 128 |
| :--- | :--- | :---: | :---: | :---: | :---: | ---: | ---: | ---: |
| Matrix multiplication | 0.002 | 0.014 | 0.12 | 0.9 | 7.2 | 56 | 460 | 3600 |
| Newton implementation | 0.13 | 0.31 | 2.3 | 13 | 95 | 700 | 5500 | 43000 |
| Naive_Mul_Padic $p$ | 0.34 | 0.42 | 1.4 | 3.3 | 8.7 | 26 | 94 | 420 |

Table 6. 256 first terms of $b=\Phi_{d}(b)$, for $p=536870923$, in milliseconds.
Although Newton iteration is faster for tiny dimensions $d \leqslant 2$, its cost growths as $d^{3}$ for larger $d$, whereas the relaxed approach only grows as $d^{2}$. For $d=1$, we notice that the number $b$ is computed with essentially one relaxed product. In the next table we report of the same computations but with the relaxed product of Section 3.2 at precision 1024; the conclusions are essentially the same:

| $d$ | 1 | 2 | 4 | 8 | 16 | 32 | 64 | 128 |
| :--- | :--- | :---: | :---: | :---: | ---: | ---: | ---: | :---: |
| Matrix multiplication | 0.014 | 0.12 | 0.98 | 7.9 | 62 | 490 | 4000 | 31000 |
| Newton implementation | 0.58 | 1.2 | 13 | 90 | 640 | 4900 | 38000 | 300000 |
| Binary_Mul_Padic $_{p}$ | 2.4 | 2.6 | 9.4 | 21 | 52 | 150 | 480 | 2300 |

Table 7. 1024 first terms of $b=\Phi_{d}(b)$, for $p=536870923$, in milliseconds.

## 6. Relaxed division

We are now to present relaxed algorithms to compute the quotient of two $p$-adic numbers. The technique is similar to power series, as treated in [Hoe02], but with subtleties.

### 6.1. Division by a "scalar"

The division of a power series in $\mathbb{K}[[x]]$ by an element of $\mathbb{K}$ is immediate, but it does not extend to $p$-adic numbers, because of the propagation of the carries. We shall introduce two new operations. Let $\beta \in M$ play the role of a "scalar". The first new operation, written mul_rem ( $a: \operatorname{Padic}_{p}, \beta: M$ ), returns the $p$-adic number $c$ with coefficients $c_{n}=$ rem $\left(\beta a_{n}, p\right)$. The second operation, written mul_quo ( $a: \operatorname{Padic}_{p}, \beta: R$ ), returns the corresponding carry, so that

$$
\begin{aligned}
\beta a & =\text { mul_rem }(a, \beta)+p \text { mul_quo }(a, \beta) \\
& =\sum_{n=0}^{\infty} \operatorname{rem}\left(\beta a_{n}, p\right) p^{n}+\sum_{n=0}^{\infty} \operatorname{quo}\left(\beta a_{n}, p\right) p^{n+1} .
\end{aligned}
$$

These operations are easy to implement, as follows:

```
class Mul_Rem_Padic_rep \({ }_{p} \triangleq\) Padic_rep \(_{p}\)
    \(a:\) Padic \(_{p}\)
    \(\beta\) : \(M\)
    constructor \(\left(\tilde{a}: \operatorname{Padic}_{p}, \tilde{\beta}: M\right)\)
        \(a:=\tilde{a}, \beta:=\tilde{\beta}\)
    method next ()
        return rem \(\left(\beta a_{n}, p\right)\)
class Mul_Quo_Padic_rep \({ }_{p} \triangleq\) Padic_rep \(_{p}\)
    \(a:\) Padic \(_{p}\)
    \(\beta: M\)
    constructor \(\left(\tilde{a}: \operatorname{Padic}_{p}, \tilde{\beta}: M\right)\)
        \(a:=\tilde{a}, \beta:=\tilde{\beta}\)
    method next ()
        return quo \(\left(\beta a_{n}, p\right)\)
```

Proposition 10. Let a be a relaxed $p$-adic number and let $\beta \in M$. If $\beta$ is invertible modulo $p$, with given inverse $\gamma=\beta^{-1} \bmod p$, then the quotient $c=a / \beta$ is recursive and c satisfies the equation

$$
c=\text { mul_rem }(a-p \text { mul_quo }(c, \beta), \gamma), \quad c_{0}=\gamma a_{0} \bmod p .
$$

If $R=\mathbb{Z}$, then $a / \beta$ can be computed up till precision $n$ using $O(n \mathrm{I}(\log p))$ bit-operations.
Proof. It is clear from the definitions that the proposed formula actually defines a recursive number. Then, from $\beta c=$ mul_rem $(c, \beta)+p$ mul_quo $(c, \beta)$, we deduce that $\beta c-p$ mul_quo $(c, \beta)=$ mul_rem $(c, \beta)$, hence

$$
c=\text { mul_rem }(\beta c-p \text { mul_quo }(c, \beta), \gamma)=\text { mul_rem }(a-p \text { mul_quo }(c, \beta), \gamma) \text {. }
$$

The functions mul_rem and mul_quo both take $O(n \mathrm{I}(\log p))$ bit-operations if $R=\mathbb{Z}$, which concludes the proof.

### 6.2. Quotient of two $\boldsymbol{p}$-adic numbers

Once the division by a "scalar" is available, we can apply a similar formula as for the division of power series of [Hoe02].

Proposition 11. Let $a$ and $b$ be two relaxed $p$-adic numbers such that $b_{0}$ is invertible of given inverse $\gamma=b_{0}^{-1} \bmod p$. The quotient $c=a / b$ is recursive and satisfies the following equation:

$$
c=\frac{a-\left(b-b_{0}\right) c}{b_{0}}, \quad c_{0}=\gamma a_{0} \bmod p
$$

If $R=\mathbb{Z}$, then $a / b$ can be computed up till precision $n$ using $O(\mathrm{I}(n \log p) \log n)$ bit-operations.

Proof. The last assertion on the cost follows from Proposition 7.
Remark 12. Notice that $p$ is not assumed to be prime, so that we can replace $p$ by $p^{k}$, and thus benefit of the monoblock strategy of Section 4.2. This does not involve a large amount of work: it suffices to write from_blocks (to_blocks $(a) /$ to_blocks $(b)$ ). Notice that this involves inverting $b_{0 \ldots p^{k}}$ modulo $p^{k}$.

### 6.3. Timings

In following Table 8 we display the computation time for our division algorithm. We compare several methods:

- The first line "Newton" corresponds to the classical Newton iteration [GG03, Algorithm 9.2] used in the zealous model.
- The second line corresponds to one call of GMP's extended g.c.d. function.
- The next two lines Naive_Mul_Padic $p_{p}$ and Binary_Mul_Padic $p_{p}$ correspond to the naive product of Section 2.6, and the relaxed one of Section 3.2.
- Then the next lines "mono Naive_Mul_Padic " and "mono Naive_Mul_Padic ${ }_{p}$ " correspond to the monoblock strategy from Section 4.2 with blocks of size 32 .
- Similarly the lines "blocks Naive_Mul_Padic ${ }_{p}$ " and "blocks Naive_Mul_Padic ${ }_{p}$ " correspond to the relaxed block strategy from Section 4.3 with blocks of size 32 .
- Finally the last line corresponds to direct computations in base $p^{32}$ (with no conversions from/to base $p$ ).

When compared to Tables 1 and 2, we observe that the cost of one division algorithm is merely that of one multiplication whenever the size becomes sufficiently large, as expected. We also observe that our "monoblock division" is faster than the zealous one for large sizes; this is even more true if we directly compute in base $p^{32}$.

| $n$ | 8 | 16 | 32 | 64 | 128 | 256 | 512 | 1024 | 2048 |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | :---: |
| Newton | 3 | 4 | 7 | 18 | 49 | 140 | 430 | 1300 | 3700 |
| GmP's extended g.c.d. $^{2}$ | 3 | 6 | 14 | 35 | 92 | 250 | 730 | 2200 | 5600 |
| Naive_Mul_Padic $_{p}$ | 4 | 7 | 15 | 35 | 95 | 300 | 1000 | 3700 | 14000 |
| Binary_Mul_Padic $_{p}$ | 9 | 21 | 44 | 93 | 200 | 420 | 920 | 2000 | 4800 |
| mono Naive_Mul_Padic $_{p}$ | 6 | 10 | 20 | 95 | 160 | 280 | 540 | 1200 | 2800 |
| mono Binary_Mul_Padic $_{p}$ | 6 | 10 | 20 | 110 | 170 | 320 | 660 | 1500 | 3600 |
| blocks Naive_Mul_Padic $_{p}$ | 10 | 16 | 27 | 70 | 190 | 430 | 900 | 1900 | 4500 |
| blocks Binary_Mul_Padic $_{p}$ | 10 | 16 | 27 | 65 | 180 | 410 | 900 | 2000 | 4500 |
| Naive_Mul_Padic_$p^{32}$ |  |  | 6 | 22 | 42 | 88 | 200 | 500 | 1500 |

Table 8. Divisions, for $p=536870923$, in microseconds.

## 7. Higher order roots

For power series in characteristic 0 , the $r$-th root $g$ of $f$ is recursive, with equation $g=$ $\int f^{\prime} /\left(r g^{r-1}\right)$ and initial condition $g_{0}=f_{0}^{1 / r}$ (see [Hoe02, Section 3.2.5] for details). This expression neither holds in small positive characteristic, nor for $p$-adic integers. In this section we propose new formulas for these two cases, which are compatible with the monoblock strategy of Section 4.2.

### 7.1. Regular case

In this subsection we treat the case when $r$ is invertible modulo $p$.
Proposition 13. Assume that $r$ is invertible modulo $p$, and let a be a relaxed invertible $p$-adic number in $R_{p}$ such that $a_{0}$ is an $r$-th power modulo $p$. Then any $r$-th root $b_{0}$ of $a_{0}$ modulo $p$ can be uniquely lifted into an $r$-th root $b$ of $a$. Moreover, $b$ is a recursive number for the equation

$$
\begin{equation*}
b=\frac{a-b^{r}+r b_{0}^{r-1} b}{r b_{0}^{r-1}} . \tag{4}
\end{equation*}
$$

The $n$ first terms of $b$ can be computed using

- $O(\log r \mathbb{M}(n) \log n)$ operations in $\mathbb{K}$, if $R=\mathbb{K}[[x]]$ and $p=x$, or
- $O(\log r \mathrm{I}(n \log p) \log n)$ bit-operations, if $R=\mathbb{Z}$.

Proof. Since $r$ is invertible modulo $p$, the polynomial $x^{r}-a$ is separable modulo $p$. Any of its roots modulo $p$ can be uniquely lifted into a root in $R_{p}$ by means of the classical Newton operator [Lan02, Proposition 7.2].

Since $a_{0}$ is invertible, so is $b_{0}$. It is therefore clear that Equation (4) uniquely defines $b$, but it is not immediately clear how to evaluate it so that it defines a recursive number. For this purpose we rewrite $b$ into $b_{0}+c$, with $c$ of valuation at least 1 :

$$
b=\frac{a-\left(b_{0}+c\right)^{r}+r b_{0}^{r-1}\left(b_{0}+c\right)}{r b_{0}^{r-1}}=\frac{a-\sum_{k=2}^{r}\binom{r}{k} b_{0}^{r-k} c^{k}+(r-1) b_{0}^{r}}{r b_{0}^{r-1}} .
$$

Since $r$ is invertible modulo $p$, we now see that it does suffice to know the terms to degree $n-1$ of $b$ in order to deduce $b_{n}$.

The latter expanded formula is suitable for an implementation but unfortunately the number of products to be performed grows linearly with $r$. Instead we modify the classical binary powering algorithm to compute the expression needed with $O(\log r)$ products only, as follows. In fact we aim at computing $\beta_{r}=\left(b_{0}+c\right)^{r}-r b_{0}^{r-1} c-b_{0}^{r}$ in a way to preserve the recursiveness. We proceed by induction on $r$.

If $r=1$, then $\beta_{r}=0$. If $r=2$ then $\beta_{2}=c^{2}$. Assume that $r=2 h$, and that $\beta_{h}$ is available by induction. From

$$
\begin{aligned}
\beta_{h}^{2} & =\left(b_{0}+c\right)^{r}+\left(h b_{0}^{h-1} c+b_{0}^{h}\right)^{2}-2\left(h b_{0}^{h-1} c+b_{0}^{h}\right)\left(\beta_{h}+h b_{0}^{h-1} c+b_{0}^{h}\right) \\
& =\left(b_{0}+c\right)^{r}-\left(h b_{0}^{h-1} c+b_{0}^{h}\right)^{2}-2\left(h b_{0}^{h-1} c+b_{0}^{h}\right) \beta_{h},
\end{aligned}
$$

we deduce that

$$
\begin{aligned}
\beta_{r} & =\beta_{h}^{2}+\left(h b_{0}^{h-1} c+b_{0}^{h}\right)^{2}+2\left(h b_{0}^{h-1} c+b_{0}^{h}\right) \beta_{h}-r b_{0}^{r-1} c-b_{0}^{r} \\
& =\beta_{h}\left(\beta_{h}+2\left(h b_{0}^{h-1} c+b_{0}^{h}\right)\right)+\left(h b_{0}^{h-1} c\right)^{2} .
\end{aligned}
$$

Since $\beta_{h}$ and $c$ have positive valuation, the recursiveness is well preserved through this intermediate expression.

On the other hand, if $r$ is odd then we can write $r=h+1$, with $h$ even, and assume that $\beta_{h}$ is available by induction. Then we have that:

$$
\begin{aligned}
\beta_{r} & =\left(b_{0}+c\right) \beta_{h}+\left(b_{0}+c\right)\left(h b_{0}^{h-1} c+b_{0}^{h}\right)-(h+1) b_{0}^{h} c-b_{0}^{h+1} \\
& =\left(b_{0}+c\right) \beta_{h}+h b_{0}^{h-1} c^{2} .
\end{aligned}
$$

Again the recursiveness is well preserved through this intermediate expression. The equation of $b$ can finally be evaluated using $O(\log r)$ products and one division. By [Hoe02, Theorem 4.1], this concludes the proof for power series. By Propositions 7 and 11 , we also obtain the desired result for $p$-adic integers.

For the computation of the $r$-th root in $\mathbb{Z} / p \mathbb{Z}$, we have implemented the algorithms of [GG03, Theorems 14.4 and 14.9]: each extraction can be done with $\tilde{O}(r \log p)$ bitoperations in average, with a randomized algorithm. This is not the bottleneck for our purpose, so we will not discuss this aspect longer in this paper.

Remark 14. Notice that ( $p$ ) is not assumed to be prime in Proposition 13. Therefore, if we actually have an $r$-th root $b$ of $a$ modulo $p^{k}$, then $b$ can be seen as a $p^{k}$-recursive number, still using Equation (4). Hence, one can directly apply the monoblock strategy of Section 4.2 to perform internal computations modulo $p^{k}$.

## 7.2. $\boldsymbol{p}$-th roots

If $\mathbb{K}$ is a field of characteristic $p$, then $f \in \mathbb{K}[[x]]$ is a $p$-th power if, and only if, $f \in$ $\mathbb{K}^{p}\left[\left[x^{p}\right]\right]$. If it exists, the $p$-th root of a power series is unique. Here, $\mathbb{K}^{p}$ represents the sub-field of the $p$-th powers of $\mathbb{K}$. By the way, let us mention that, for a general effective field $\mathbb{K}$, Fröhlich and Shepherdson have shown that testing if an element is a $p$-th power is not decidable [FS56, Section 7] (see also the example in [Gat84, Remark 5.10]).

In general, for $p$-adic numbers, an $r$-th root extraction can be almost as complicated as the factorization of a general polynomial in $R_{p}[x]$. For instance, with $R=\mathbb{Z}[\sqrt{2}]$ and $p=\sqrt{2}$ we have that $r=2=p^{2}$ has valuation 2 in $R_{p}$. We will not cover such a general situation. We will only consider the case of the $p$-adic integers, that is for when $R=\mathbb{Z}$ and $p$ is prime.

From now on, let $a$ denote a $p$-adic integer in $\mathbb{Z}_{p}$ from which we want to extract the $p$-th root (if it exists). If the valuation of $a$ is not a multiple of $p$, then $a$ is not a $p$-th power. If it is a multiple of $p$, then we can factor out $p^{\mathrm{val} a}$ and assume that $a$ has valuation 0 . The following lemma is based on classical techniques, we briefly recall its proof for completeness:

Lemma 15. Assume that $p$ is prime, and let $a \in \mathbb{Z}_{p}$ be invertible.

- If $p \geqslant 3$, then a is a p-th power if, and only if, $a_{0}+p a_{1}=a_{0}^{p}$ modulo $p^{2}$. In this case there exists only one $p$-th root.
- If $p=2$, then $a$ is a p-th power if, and only if, $a_{1}=a_{2}=0$. In this case there exist exactly two square roots.

Proof. If $a=b^{p}$ in $\mathbb{Z}_{p}$ then $b_{0}=a_{0}$. After the translation $x=a_{0}+y$ in $x^{p}-a=0$, we focus on the equation $\left(b_{0}+y\right)^{p}-a=0$, which expands to

$$
\begin{equation*}
h(y)=y^{p}+\sum_{i=1}^{p-1}\binom{p}{i} b_{0}^{p-i} y^{i}-\left(a-b_{0}^{p}\right)=0 . \tag{5}
\end{equation*}
$$

For any $i \in\{1, \ldots, p-1\}$, the coefficient $\binom{p}{i}$ has valuation at least one because $p$ is prime. Reducing the latter equation modulo $p^{2}$, it is thus necessary that $a_{0}+p a_{1}=b_{0}^{p}$ modulo $p^{2}$.

Assume now that $a_{0}+p a_{1}=b_{0}^{p}$ holds modulo $p^{2}$. After the change of variables $y$ by $p z$ and division by $p^{2}$, we obtain

$$
\begin{equation*}
\tilde{h}(z)=p^{p-2} z^{p}+\sum_{i=2}^{p-1}\binom{p}{i} b_{0}^{p-i} p^{i-2} z^{i}+b_{0}^{p-1} z-\frac{a-b_{0}^{p}}{p^{2}}=0 . \tag{6}
\end{equation*}
$$

We distinguish two cases: $p \geqslant 3$ and $p=2$.
If $p \geqslant 3$, then any root of $\tilde{h}$ must be congruent to $\left(a-b_{0}^{p}\right) /\left(b_{0}^{p-1} p^{2}\right)$ modulo $p$. Since $\tilde{h}^{\prime}\left(\left(a-b_{0}^{p}\right) /\left(b_{0}^{p-1} p^{2}\right)\right)=b_{0}^{p-1}$ has valuation 0 , the Newton operator again ensures that $\tilde{h}$ has exactly one root [Lan02, Proposition 7.2].

If $p=2$, then $\tilde{h}(z)$ rewrites into $z^{2}+b_{0} z-\frac{a-b_{0}^{2}}{p^{2}}=z^{2}+z-\frac{a-1}{4}$. Since $\tilde{h}^{\prime}(z)=$ $b_{0} \bmod p=1 \bmod 2$, any root of $\tilde{h}$ modulo 2 can be lifted into a unique root of $\tilde{h}$ in $\mathbb{Z}_{2}$. The possible roots being 0 and 1 , this gives the extra condition $a_{2}=0$.

### 7.3. Square roots in base 2

In the following proposition we show that the square root of a 2 -adic integer can be expressed into a recursive number that can be computed with essentially one relaxed product.

Proposition 16. Let a be a relaxed 2-adic integer in $\mathbb{Z}_{2}$ with $a_{0}=1$ and $a_{1}=a_{2}=0$. Let $b$ be a square root of $a$, with $b_{0}=1$ and $b_{1}$ being 0 or 1 , and let $c=\left(b-b_{0}-2 b_{1}\right) / 4$, and $\tilde{a}=\left(a-\left(b_{0}+2 b_{1}\right)^{2}\right) / 8$. Then $c$ is a recursive number with initial condition $c_{0}=\tilde{a}_{0}$ and equation

$$
\begin{equation*}
c=\frac{\tilde{a}-2 c^{2}}{b_{0}+2 b_{1}} . \tag{7}
\end{equation*}
$$

The $n$ first terms of $b$ can be computed using $O(1(n \log p) \log n)$ bit-operations.
Proof. Equation (7) simply follows from

$$
\left(b_{0}+2 b_{1}+4 c\right)^{2}-a=\left(b_{0}+2 b_{1}\right)^{2}+8\left(b_{0}+2 b_{1}\right) c+16 c^{2}-a=0 .
$$

The cost is a consequence of Propositions 7 and 11.
Remark 17. As in the preceding regular case, we can see $c$ as a $p^{k}$-recursive number as soon as $c$ is known modulo $p^{k}$. In fact letting $C=C_{0}+\tilde{C}$, with $C_{0}=C \bmod p^{k}$, Equation (7) rewrites into

$$
\begin{aligned}
\left(b_{0}+2 b_{1}\right) C & =\tilde{a}-2\left(C_{0}+\left(C-C_{0}\right)\right)^{2} \\
& =\tilde{a}-2 C_{0}^{2}-4 C_{0}\left(C-C_{0}\right)-2\left(C-C_{0}\right)^{2},
\end{aligned}
$$

which gives

$$
\left(b_{0}+2 b_{1}+4 C_{0}\right) C=\tilde{a}+2 C_{0}^{2}-2\left(C-C_{0}\right)^{2} .
$$

The latter equation implies that $C$ is $p^{k}$-recursive, so that we can naturally benefit of the monoblock strategy from Section 4.2.

## 7.4. $p$-th roots in base $p$

In this subsection we assume that $p$ is an odd prime integer. We will show that the $p$-th root is recursive and can be computed using similar but slightly more complicated formulas than in the regular case.

Proposition 18. Let $a$ be an invertible relaxed $p$-adic integer in $\mathbb{Z}_{p}$ such that $a=$ $a_{0}^{p} \bmod p^{2}$. Let $b$ denote the $p$-th root of $a$, with $b_{0}=a_{0}$, let $\tilde{a}=\left(a-\left(b_{0}+p b_{1}\right)^{p}\right) / p^{2}$, and let $c=\left(b-b_{0}\right) / p$. Then $c$ is a recursive number with initial condition $c_{0}=b_{1}$ and equation

$$
\begin{equation*}
c=c_{0}+\frac{\tilde{a}-\gamma_{p}}{\left(b_{0}+p c_{0}\right)^{p-1}}, \tag{8}
\end{equation*}
$$

where

$$
\gamma_{r}=\frac{\left(b_{0}+p c\right)^{r}-r p\left(b_{0}+p c_{0}\right)^{r-1}\left(c-c_{0}\right)-\left(b_{0}+p c_{0}\right)^{r}}{p^{2}}, \text { for all } r \geqslant 0 .
$$

The $n$ first terms of $b$ can be computed with $O(\log p \mathbf{I}(n \log p) \log n)$ bit-operations.
Proof. As a shorthand we let $\beta=b_{0}+p c_{0}$ and $d=c-c_{0}$. Equation (8) simply follows from

$$
\begin{aligned}
b^{p}-a & =(\beta+p d)^{p}-a \\
& =p^{2} \gamma_{p}+p^{2} \beta^{p-1} d+\beta^{p}-a \\
& =p^{2} \gamma_{p}+p^{2} \beta^{p-1} d-p^{2} \tilde{a}=0 .
\end{aligned}
$$

Since

$$
\gamma_{p}=\sum_{i=2}^{p}\binom{p}{i} \beta^{p-i} p^{i-2} d^{i},
$$

and since $d$ has positive valuation, Equation (8) actually defines $c$ as a recursive number.
As in the regular case, we need to provide an efficient way to compute $\gamma_{r}$. We proceed by induction on $r$. If $r=1$, then $\gamma_{r}=0$. If $r=2$, then $\gamma_{r}=d^{2}$, which preserves the recursiveness. Assume now that $r=2 h$ and that $\gamma_{h}$ is available by induction. From

$$
\begin{aligned}
\left(p^{2} \gamma_{h}\right)^{2}= & (\beta+p d)^{r}+\left(h p \beta^{h-1} d+\beta^{h}\right)^{2} \\
& -2\left(h p \beta^{h-1} d+\beta^{h}\right)\left(p^{2} \gamma_{h}+h p \beta^{h-1} d+\beta^{h}\right) \\
= & (\beta+p d)^{r}-\left(h p \beta^{h-1} d+\beta^{h}\right)^{2}-2\left(h p \beta^{h-1} d+\beta^{h}\right) p^{2} \gamma_{h},
\end{aligned}
$$

we deduce that

$$
\begin{aligned}
p^{2} \gamma_{r} & =\left(p^{2} \gamma_{h}\right)^{2}+\left(h p \beta^{h-1} d+\beta^{h}\right)^{2}+2\left(h p \beta^{h-1} d+\beta^{h}\right) p^{2} \gamma_{h}-r p \beta^{r-1} d-\beta^{r} \\
& =p^{2} \gamma_{h}\left(p^{2} \gamma_{h}+2\left(h p \beta^{h-1} d+\beta^{h}\right)\right)+\left(h p \beta^{h-1} d\right)^{2},
\end{aligned}
$$

whence

$$
\gamma_{r}=\gamma_{h}\left(p^{2} \gamma_{h}+2\left(h p \beta^{h-1} d+\beta^{h}\right)\right)+\left(h \beta^{h-1} d\right)^{2} .
$$

Since $\gamma_{h}$ and $d$ have positive valuation, the recursiveness is well preserved through this intermediate expression.

On the other hand, if $r$ is odd, then we can write $r=h+1$, with $h$ even, and assume that $\gamma_{h}$ is available by induction. Then we have that:

$$
\begin{aligned}
p^{2} \gamma_{r} & =p^{2}(\beta+p d) \gamma_{h}+(\beta+p d)\left(h p \beta^{h-1} d+\beta^{h}\right)-(h+1) p \beta^{h} d-\beta^{h+1} \\
& =p^{2}(\beta+p d) \gamma_{h}+h p^{2} \beta^{h-1} d^{2},
\end{aligned}
$$

whence

$$
\gamma_{r}=(\beta+p d) \gamma_{h}+h \beta^{h-1} d^{2}
$$

which again preserves the recursiveness. Finally the equation of $d$ can be evaluated with $O(\log r)$ products and one division, which concludes the proof by Propositions 7 and 11.

REMARK 19. As in the regular case, we can see $c$ as a $p^{k}$-recursive number as soon as $c$ is known modulo $p^{k}$. In fact, letting $C=C_{0}+D$, with $C_{0}=C \bmod p^{k}$, Equation (8) rewrites into

$$
C=C_{0}+\frac{\tilde{A}-\Gamma_{p}}{\left(b_{0}+p C_{0}\right)^{p-1}},
$$

where $\tilde{A}=\left(a-\left(b_{0}+p C_{0}\right)^{p}\right) / p^{2}$, and

$$
\Gamma_{r}=\frac{\left(b_{0}+p C\right)^{r}-r p\left(b_{0}+p C_{0}\right)^{r-1}\left(C-C_{0}\right)-\left(b_{0}+p C_{0}\right)^{r}}{p^{2}}, \text { for all } r \geqslant 0
$$

Notice that division by $p^{2}$ in base $p^{k}$, with $k>2$, is equivalent to multiplication by $p^{k-2}$ and division by $p^{k}$, which corresponds to a simple shift. Then $\Gamma_{p}$ can be computed by recurrence with the same formula as $\gamma_{p}$, mutatis mutandis. In this way $C$ is $p^{k}$-recursive, so that we can naturally benefit of the monoblock strategy of Section 4.2.

### 7.5. Timings

In the following table, we give the computation time of the square root using our fast relaxed product of Section 4.3 that has been reported in Table 5. Since, in terms of performances, the situation is very similar to the division, we only compare to the zealous implementation in PARI/GP version 2.3.5.

| $n$ | 8 | 16 | 32 | 64 | 128 | 256 | 512 | 1024 | 2048 |
| :--- | ---: | :--- | :--- | :--- | :--- | :--- | :---: | :---: | :---: |
| blocks Binary_Mul_Padic $_{p}$ | 14 | 22 | 39 | 91 | 230 | 520 | 1100 | 2400 | 5400 |
| PARI/GP | 5 | 11 | 28 | 74 | 210 | 670 | 2300 | 8100 | 30000 |

Table 9. Square root, for $p=536870923$, in microseconds.

## 8. Conclusion

From more than a decade a major stream in complexity theory for computer algebra has spread the idea that high level algorithms must be parameterized in terms of a small number of elementary operations (essentially integer, polynomial and matrix multiplication), so that the main goal in algorithm design consists in reducing as fast as possible to these operations. Although many asymptotically efficient algorithms have been developed along these lines, an overly doctrinaire application of this philosophy tends to be counterproductive.

For example, when it comes to computations in completions, we have seen that there are two general approaches: Newton's method and the relaxed (or lazy) approach. It is often believed that Newton's method is simply the best, because it asymptotically leads to the same complexity as integer or polynomial multiplication. However, this "reduction" does not take into account possible sparsity in the data, non asymptotic input sizes, more involved types of equations (such as partial differential equations), etc. In the literature, these difficulties are often evacuated by picking problems, not according to their practical relevance, but rather to their aptitude to fit in an idealized world where standard complexity arguments can be applied.

In this paper, we have demonstrated that, in the area of computations with $p$-adic numbers, the relaxed approach is often more efficient than methods based on Newton iteration. The gains are sometimes spectacular: in Tables 6 and 7, we have shown that Hensel lifting in high dimensions can become more than 100 times faster, when using the relaxed approach. At other moments, we were ourselves surprised: in Table 8, we see that, even for the division of $p$-adic numbers, a naive implementation of the relaxed product yields better performances than a straightforward use of GMP, whenever $p$ is sufficiently large.

Of course, the detailed analysis of the mutual benefits of both approaches remains an interesting subject. On the one hand, Newton iteration can be improved using blockwise techniques [Ber00, Hoe10]. On the other hand, the relaxed implementation can be improved for small sizes by ensuring a better transition between hardware and long integers, and massive inlining. At huge precisions, the recursive blockwise technique from [Hoe07a] should also become useful. Finally, "FFT-caching" could still be used in a more systematic way, and in particular for the computation of squares.

To conclude our comparison between Newton iteration and the relaxed approach, we would like to stress that, under most favourable circumstances, Newton iteration can only be hoped to be a small constant times faster than the relaxed approach, since the overhead $O(\log n)$ of relaxed multiplication should really be read as $(1 / 2) \log (n / 128)$ or less. In other extreme cases, Newton iteration is several hundreds times slower, or even does not apply at all (e.g. for the resolution of partial differential equations).

Let us finally mention that the relaxed resolution of recursive systems of equations has been extended to more general systems of implicit equations [Hoe09]. The computation of such local solutions is the central task of the polynomial system solver called Kronecker (see [DL08] for an introduction). We are confident that the results of [Hoe09], which were presented in the power series context, extend to more general completions, and that the relaxed model will lead to an important speed-up.

## Bibliography

[Ber00] D. Bernstein. Removing redundancy in high precision Newton iteration. Available from http://cr.yp.to/fastnewton.html, 2000.
[BK78] R. P. Brent and H. T. Kung. Fast algorithms for manipulating formal power series. Journal of the ACM, 25:581-595, 1978.
[Bre76] R. P. Brent. The complexity of multiprecision arithmetic. In R. S. Anderssen and R. P. Brent, editors, Complexity of computational problem solving, pages 126-165, Brisbane, 1976. University of Queensland Press.
[CC90] D. V. Chudnovsky and G. V. Chudnovsky. Computer algebra in the service of mathematical physics and number theory (computers in mathematics, stanford, ca, 1986). In Lect. Notes in Pure and Applied Math., volume 125, pages 109-232, New-York, 1990. Dekker.
[CK91] D. G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. Acta Informatica, 28:693-701, 1991.
[DL08] C. Durvye and G. Lecerf. A concise proof of the Kronecker polynomial system solver from scratch. Expositiones Mathematicae, 26(2):101-139, 2008.
[DS04] S. De Smedt. p-adic arithmetic. The Mathematica Journal, 9(2):349-357, 2004.
[FS56] A. Fröhlich and J. C. Shepherdson. Effective procedures in field theory. Philos. Trans. Roy. Soc. London. Ser. A., 248:407-432, 1956.
[Für07] M. Fürer. Faster integer multiplication. In Proceedings of the Thirty-Ninth ACM Symposium on Theory of Computing (STOC 2007), pages 57-66, San Diego, California, 2007.
$[\mathbf{G}+\mathbf{9 1}]$ T. Granlund et al. GMP, the GNU multiple precision arithmetic library. Available from http://www.swox.com/gmp, 1991.
[Gat84] J. von zur Gathen. Hensel and Newton methods in valuation rings. Math. Comp., 42(166):637-661, 1984.
[GG03] J. von zur Gathen and J. Gerhard. Modern computer algebra. Cambridge University Press, Cambridge, second edition, 2003.
[H+02] J. van der Hoeven et al. Mathemagix. Available from http://www.mathemagix.org, 2002.
[HH09] W. B. Hart and D. Harvey. FLINT 1.5.1: Fast library for number theory. Available from http://www.flintlib.org, 2009.
[Hoe97] J. van der Hoeven. Lazy multiplication of formal power series. In W. W. Küchlin, editor, Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation (ISSAC 1997), pages 17-20, Maui, Hawaii, July 1997.
[Hoe99] J. van der Hoeven. Fast evaluation of holonomic functions. Theoretical Computer Science, 210:199-215, 1999.
[Hoe01] J. van der Hoeven. Fast evaluation of holonomic functions near and in singularities. J. Symbolic Comput., 31:717-743, 2001.
[Hoe02] J. van der Hoeven. Relax, but don't be too lazy. J. Symbolic Comput., 34(6):479-542, 2002.
[Hoe07a] J. van der Hoeven. New algorithms for relaxed multiplication. J. Symbolic Comput., 42(8):792-802, 2007.
[Hoe07b] Joris van der Hoeven. Efficient accelero-summation of holonomic functions. J. Symbolic Comput., 42(4):389-428, 2007.
[Hoe09] J. van der Hoeven. Relaxed resolution of implicit equations. Technical report, HAL, 2009. http://hal.archives-ouvertes.fr/hal-00441977/fr/.
[Hoe10] J. van der Hoeven. Newton's method and FFT trading. J. Symbolic Comput., 45(8):857878, 2010.
[Kat07] S. Katok. p-adic analysis compared with real, volume 37 of Student Mathematical Library. American Mathematical Society, Providence, RI, 2007.
[Kob84] N. Koblitz. p-adic numbers, p-adic analysis, and zeta-functions, volume 58 of Graduate Texts in Mathematics. Springer-Verlag, New York, second edition, 1984.
[Lan02] S. Lang. Algebra, volume 211 of Graduate Texts in Mathematics. Springer-Verlag, third edition, 2002.
[PAR08] The PARI Group, Bordeaux. PARI/GP, version 2.3.5, 2008. Available from http://pari.math.u-bordeaux.fr/.
[S+09] W. A. Stein et al. Sage Mathematics Software (Version 4.2.1). The Sage Development Team, 2009. Available from http://www.sagemath.org.
[SS71] A. Schönhage and V. Strassen. Schnelle Multiplikation grosser Zahlen. Computing, 7:281292, 1971.
[Wan84] P. S. Wang. Implementation of a $p$-adic package for polynomial factorization and other related operations. In EUROSAM 84 (Cambridge, 1984), volume 174 of Lecture Notes in Comput. Sci., pages 86-99. Springer, Berlin, 1984.


[^0]:    *. This work has been partly supported by the French ANR-09-JCJC-0098-01 MaGIX project, and by the Digiteo 2009-36HD grant of the Région Ile-de-France.

